



NetApp storage deployment procedure (part 1)

FlexPod

NetApp
June 03, 2021

Table of Contents

- NetApp storage deployment procedure (part 1) 1
- NetApp storage controller AFF2xx series installation..... 1
- NetApp ONTAP 9.4 1
- Continuation of Node A configuration and cluster configuration 4
- Continuation of storage cluster configuration..... 9

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF2xx series installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

1. Access the [HWU](#) application to view the system configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFF2XX Series prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections: Electrical Requirements, Supported Power Cords, and Onboard Ports and Cables.

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A220 Documentation](#).

NetApp ONTAP 9.4

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.4 Software Setup Guide](#).



This system is set up in a two-node switchless cluster configuration.

The following table shows ONTAP 9.4 installation and configuration information.

Cluster detail	Cluster detail value
Cluster node A IP address	<<var_nodeA_mgmt_ip>>
Cluster node A netmask	<<var_nodeA_mgmt_mask>>
Cluster node A gateway	<<var_nodeA_mgmt_gateway>>
Cluster node A name	<<var_nodeA>>
Cluster node B IP address	<<var_nodeB_mgmt_ip>>
Cluster node B netmask	<<var_nodeB_mgmt_mask>>

Cluster detail	Cluster detail value
Cluster node B gateway	<<var_nodeB_mgmt_gateway>>
Cluster node B name	<<var_nodeB>>
ONTAP 9.4 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<<var_dns_server_ip>>
NTP server IP (you can enter more than one)	<<var_ntp_server_ip>>

Configure Node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot.

```
autoboot
```

3. Press Ctrl-C to enter the Boot menu.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

17. While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and `y` to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter `y` to perform an upgrade.

6. Select e0M for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation of Node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.4 boots on the node for the first time.



The node and cluster setup procedure has changed slightly in ONTAP 9.4. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up Node A.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
    "help" or "?" - if you want to have a question clarified,
    "back" - if you want to change previously answered questions, and
    "exit" or "quit" - if you want to quit the cluster setup wizard.
    Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]:
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>
Enter the node management interface default gateway:
<<var_nodeA_mgmt_gateway>>
A node management interface on port e0M with IP address
<<var_nodeA_mgmt_ip>> has been created.
Use your web browser to complete cluster setup by accessing
https://<<var\_nodeA\_mgmt\_ip>>
Otherwise, press Enter to complete cluster setup using the command line
interface:
```

2. Navigate to the IP address of the node's management interface.

Cluster setup can also be performed by using the CLI. This document describes cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup to configure the cluster.
4. Enter <<var_clustertype>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster Name

Nodes

i Not sure all nodes have been discovered? [Refresh](#)

FAS2650 62165000092 <input type="text"/>	HA PAIR	FAS2650 62165000093 <input type="text"/>
--	---------	--

Cluster Configuration: Switched Cluster Switchless Cluster

? Username admin

Password

Confirm Password

Cluster Base License (Optional)

i For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional)

i Cluster Base License is mandatory to add Feature Licenses.

Submit

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter `<<var_clustermgmt_ip>>` in the Cluster Management IP Address field, `<<var_clustermgmt_mask>>` in the Netmask field, and `<<var_clustermgmt_gateway>>` in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter `<<var_nodeA_mgmt_ip>>` for node B.

- d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter <<var_ntp_server_ip>> in the Primary NTP Server field.

You can also enter an alternate NTP server.

8. Configure the support information.

- a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
- b. Enter the SMTP mail host and email address for event notifications.

You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:

<input checked="" type="checkbox"/>	Email	SMTP Mail Host <input type="text"/>	Email Addresses <input type="text" value="Separate email addresses with a comma..."/>
-------------------------------------	--------------	---	---

<input type="checkbox"/>	SNMP	SNMP Trap Host <input type="text"/>
--------------------------	-------------	---

<input type="checkbox"/>	Syslog	Syslog Server <input type="text"/>
--------------------------	---------------	--

Submit

9. When indicated that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set on-board UTA2 ports personality

1. Verify the current mode and the current type of the ports by running the `ucadmin show` command.

```
AFF A220::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
-----	-----	-----	-----	-----	-----	-----

AFF A220_A	0c	fc	target	-	-	online
AFF A220_A	0d	fc	target	-	-	online
AFF A220_A	0e	fc	target	-	-	online
AFF A220_A	0f	fc	target	-	-	online
AFF A220_B	0c	fc	target	-	-	online
AFF A220_B	0d	fc	target	-	-	online
AFF A220_B	0e	fc	target	-	-	online
AFF A220_B	0f	fc	target	-	-	online

8 entries were displayed.

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by using the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode  
cna -type target
```

The ports must be offline to run the previous command. To take a port offline, run the following command:

```
`network fcp adapter modify -node <home node of the port> -adapter <port  
name> -state down`
```



If you changed the port personality, you must reboot each node for the change to take effect.

Rename management logical interfaces (LIFs)

To rename the management LIFs, complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_node_mgmt_lif_AFF A220_B_1 -newname AFF A220-02_mgmt1
```

Set auto-revert on cluster management

Set the `auto-revert` parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-  
revert true
```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>>  
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>  
system service-processor network modify -node <<var_nodeB>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>>  
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.

This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.



Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true  
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```

The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node
<<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node
<<var_nodeB>>
```

Create jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove data ports from default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```

net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y

```

Configure IFGRP LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Make sure the switch is properly configured.

From the cluster prompt, complete the following steps.

```

ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d
ifgrp create -node << var_nodeB>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d

```

Configure jumbo frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (that usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF A220::> network port modify -node node_A -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>:a0a-<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-
<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_A_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_B_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<mgmt_vlan_id>>

```

Create aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks; you can add disks to an aggregate when additional storage is required.

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1`_`nodeA` is online.

Configure time zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is `America/New York`. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>  
snmp location "<<var_snmp_location>>"  
snmp init 1  
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select `des` as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
`nfs create -vserver Infra-SVM -udp disabled`
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
`vserver nfs modify -vserver Infra-SVM -vstorage enabled`  
`vserver nfs show`
```



Commands are prefaced by `vserver` in the command line because storage virtual machines were previously called servers.

Configure NFSv3 in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
ESXi host A NFS IP address	<<var_esxi_hostA_nfs_ip>>
ESXi host B NFS IP address	<<var_esxi_hostB_nfs_ip>>

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 2 -protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS C-Series servers are added.

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate
aggr1_nodeA -size 1GB -type DP
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate
aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m01 -type LS -schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol
snapmirror show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.

Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name  
Infra-SVM -ca Infra-SVM -type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.netapp.com  
-type server -size 2048 -country US -state "North Carolina" -locality  
"RTP" -organization "NetApp" -unit "FlexPod" -email-addr  
"abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256  
-vserver Infra-SVM
```

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...  
Example: security ssl modify -vserver Infra-SVM -server-enabled true  
-client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common  
-name infra-svm.netapp.com
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be
    interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http
-vserver <<var_clustername>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled
true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_nodeA -size 500GB -state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Enable deduplication in ONTAP

To enable deduplication on appropriate volumes, run the following commands:

```
volume efficiency on -vserver Infra-SVM -volume infra_datastore_1
volume efficiency on -vserver Infra-SVM -volume esxi_boot
```

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

```

lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware -space-reserve disabled

```



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail value
Storage node A iSCSI LIF01A	<<var_nodeA_iscsi_lif01a_ip>>
Storage node A iSCSI LIF01A network mask	<<var_nodeA_iscsi_lif01a_mask>>
Storage node A iSCSI LIF01B	<<var_nodeA_iscsi_lif01b_ip>>
Storage node A iSCSI LIF01B network mask	<<var_nodeA_iscsi_lif01b_mask>>
Storage node B iSCSI LIF01A	<<var_nodeB_iscsi_lif01a_ip>>
Storage node B iSCSI LIF01A network mask	<<var_nodeB_iscsi_lif01a_mask>>
Storage node B iSCSI LIF01B	<<var_nodeB_iscsi_lif01b_ip>>
Storage node B iSCSI LIF01B network mask	<<var_nodeB_iscsi_lif01b_mask>>

1. Create four iSCSI LIFs, two on each node.

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface show

```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

Detail	Detail Value
Storage node A NFS LIF 01 IP	<<var_nodeA_nfs_lif_01_ip>>
Storage node A NFS LIF 01 network mask	<<var_nodeA_nfs_lif_01_mask>>
Storage node B NFS LIF 02 IP	<<var_nodeB_nfs_lif_02_ip>>
Storage node B NFS LIF 02 network mask	<<var_nodeB_nfs_lif_02_mask>>

1. Create an NFS LIF.

```

network interface create -vserver Infra-SVM -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_ip>> -netmask <<
var_nodeA_nfs_lif_01_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_ip>> -netmask <<
var_nodeB_nfs_lif_02_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface show

```

Add infrastructure SVM administrator

The following table lists the information needed to complete this configuration.

Detail	Detail Value
Vsmgmt IP	<<var_svm_mgmt_ip>>
Vsmgmt network mask	<<var_svm_mgmt_mask>>
Vsmgmt default gateway	<<var_svm_mgmt_gateway>>

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```

network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true

```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```

network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>>
network route show

```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver Infra-SVM
```

[Next: Cisco UCS C-Series Rack Server Deployment Procedure](#)

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.