




Clustered Data ONTAP® 8.2

SAN Express Setup Guide for 62xx Systems



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1(408) 822-6000
Fax: +1(408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-07998_A0
May 2013

Contents

Deciding whether to use this guide	5
Required network information	6
Unpacking and mounting components	10
Locating the disk shelves for each controller	10
Component mounting positions	11
Cabling the 62xx system	13
Labeling the cables	14
Cabling the disk shelf alternate control path ports	14
Cabling the disk shelf SAS data ports	15
Cabling the HA pair NVRAM ports	22
Cabling the console ports	22
Cabling the management ports	23
Cabling the private cluster interconnect for NetApp CN1610 switches	24
Cabling the private cluster ISLs for NetApp CN1610 switches	26
Cabling the Ethernet data networks	26
Cabling the Fibre Channel data fabrics	27
Cabling the power supplies	28
Setting up the cluster	29
Creating the cluster on the first node	30
Joining a node to the cluster	31
Troubleshooting boot failure of a node	32
Configuring cluster high availability in a two-node cluster	33
Synchronizing the system time across the cluster	34
Reducing root volume size to prevent warning messages	35
Configuring AutoSupport	37
Configuring a node for troubleshooting and maintenance	38
Verifying cabling	40
Downloading and installing System Manager	41
Provisioning test LUNs for verifying cluster setup	42
Creating the test aggregate	42
Creating the test Vserver	43
Creating the test LUNs	44

Setting up test hosts and verifying system operation	45
Verifying your host configuration	46
Installing the host utilities	47
Installing or enabling MPIO	47
Connecting hosts to targets	48
Discovering LUNs and formatting disks	48
Verifying setup by writing and reading data	49
Testing storage failover	49
Deleting the test provisioning	50
Where to go next	52
Copyright information	53
Trademark information	54
How to send your comments	55
Index	56

Deciding whether to use this guide

This guide describes how to complete the initial mounting, cabling, and typical configuration of a SAN configuration using clustered Data ONTAP 8.2 software and 62xx controllers. Procedures are based on the OnCommand System Manager GUI. The guide covers commonly used options and provides minimal background information. The examples in this guide depict a six-node cluster, but you can also use the guide to configure two-node, four-node, or eight-node clusters.

This guide is intended for customers who want to install a storage system cluster themselves. It starts with the components delivered to your site and ends with the system tested and ready for storage provisioning. You should use this guide if you want a standard configuration following NetApp best practices, and you do not want a lot of conceptual background information for the tasks.

- This guide is for switched clusters using NetApp cluster switches. It does not cover two-node switchless clusters or single-node clusters.
- This guide assumes that all nodes in the cluster are the same hardware model and use SAS-connected disk shelves.
- This guide assumes a new installation and not an upgrade from a previous version of Data ONTAP software.
- This guide assumes that you download and run OnCommand System Manager 2.2 or later for all applicable tasks.

This guide does not include procedures using the Data ONTAP CLI except when the CLI is the only way to complete a task.

If these assumptions are not correct for your installation, or if you want more conceptual background information, you should refer to the following documentation instead:

- *Installation and Setup Instructions* for your hardware platform model
- *Universal SAS and ACP Cabling Guide*
- *Data ONTAP Software Setup Guide for Cluster-Mode* (for new systems)
- *Clustered Data ONTAP Upgrade and Revert/Downgrade Guide* (for upgrades)
- *Clustered Data ONTAP SAN Administration Guide*
- *Clustered Data ONTAP Switch Setup Guide for Cisco Switches*
- *OnCommand System Manager Online Help* (available both from within the product and as a PDF download)

This documentation is available from the Product Documentation section of the NetApp Support Site.

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Required network information

You must get IP addresses and other network information from your network administrator before you begin your SAN configuration.

Management IP addresses and host names

You need a management IP address and host name for each node and each cluster switch.

Type of information		Your values
Cluster	Cluster name	
	Home Ethernet port	
	IP address	
	Net mask	
Node 1	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 2	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 3	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 4	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	

Type of information		Your values
Node 5	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 6	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 7	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Node 8	Host name	
	Ethernet port	e0M
	IP address	
	Net mask	
Cluster switch 1	Host name	
	IP address	
	Net mask	
Cluster switch 2	Host name	
	IP address	
	Net mask	

iSCSI IP addresses

If you plan to use iSCSI for the test configuration, you need IP addresses for the Ethernet ports used for iSCSI. You should have at least two ports per node for redundancy.

These addresses are for a single test virtual storage server (Vserver). Each production Vserver using iSCSI has its own set of IP addresses.

iSCSI port		IP address	Net mask
Node 1	Port e__ __		
Node 1	Port e__ __		
Node 1 (optional)	Port e__ __		
Node 1 (optional)	Port e__ __		
Node 2	Port e__ __		
Node 2	Port e__ __		
Node 2 (optional)	Port e__ __		
Node 2 (optional)	Port e__ __		
Node 3	Port e__ __		
Node 3	Port e__ __		
Node 3 (optional)	Port e__ __		
Node 3 (optional)	Port e__ __		
Node 4	Port e__ __		
Node 4	Port e__ __		
Node 4 (optional)	Port e__ __		
Node 4 (optional)	Port e__ __		
Node 5	Port e__ __		
Node 5	Port e__ __		
Node 5 (optional)	Port e__ __		
Node 5 (optional)	Port e__ __		
Node 6	Port e__ __		
Node 6	Port e__ __		
Node 6 (optional)	Port e__ __		
Node 6 (optional)	Port e__ __		
Node 7	Port e__ __		
Node 7	Port e__ __		
Node 7 (optional)	Port e__ __		

iSCSI port		IP address	Net mask
Node 7 (optional)	Port e__ __		
Node 8	Port e__ __		
Node 8	Port e__ __		
Node 8 (optional)	Port e__ __		
Node 8 (optional)	Port e__ __		

Unpacking and mounting components

As you unpack and mount the cluster components, it is important to track which disk shelves go with which storage controllers. Keeping track of this information and locating the disk shelves with their storage controllers simplifies cabling and ensures that each controller has its root volume and Data ONTAP software.

The factory labels disk shelves with the storage controller to which they should connect.

Note: If you ordered the cluster in NetApp cabinets, the components are already mounted. The disk shelf containing the root volume is mounted next to its storage controller.

Locating the disk shelves for each controller

Before mounting components, you must locate the disk shelves that belong to each controller. The factory installs the root volume and Data ONTAP software on one disk shelf for each controller. By cabling the correct disk shelves to each controller, you ensure that each controller has its root volume.

Before you begin

The components must not be in the rack or cabinet so that you can see the labels on the top of the component chassis.

About this task

Each disk shelf has a sticker on the top of its chassis that lists the serial number of the controller to which the shelf connects. The sticker also lists the stack number for configurations with more than one stack of disk shelves per controller.


Steps

1. As you unpack the disk shelves and storage controllers, check the labels on the top of each chassis.

The SYSTEM SN is the identifier of the storage controller and the Shelf Chassis SN is the identifier of the disk shelf.

2. Using the label information, group or mark the disk shelves and controllers so that you can cable the disk shelves to the correct storage controllers after the components are mounted.

The following illustration is an example of the label for disk shelves.



Component mounting positions

Mounting the components in the rack or cabinet in the recommended positions simplifies cabling and service access.

Depending on the number of disk shelves and size of the rack or cabinet, you might need multiple racks or cabinets. The general rule is to keep the components of an HA pair together to simplify cabling.

The following table shows a typical mounting arrangement with two disk shelves per node. A 62xx system requires two typical racks or cabinets.

Each disk shelf has a label on the top of its chassis that lists the storage controller to which that disk shelf connects. Note that the label is not visible after mounting.

Rack/cabinet 1	Rack/cabinet 2	Rack/cabinet 3
Cluster switch 1		
Cluster switch 2		
Disk shelf 2 for Node 1	Disk shelf 2 for Node 3	Disk shelf 2 for Node 5
Disk shelf 1 for Node 1	Disk shelf 1 for Node 3	Disk shelf 1 for Node 5
Node 1 controller	Node 3 controller	Node 5 controller
Node 2 controller	Node 4 controller	Node 6 controller
Disk shelf 1 for Node 2	Disk shelf 1 for Node 4	Disk shelf 1 for Node 6

12 | SAN Express Setup Guide for 62xx Systems

Rack/cabinet 1	Rack/cabinet 2	Rack/cabinet 3
Disk shelf 2 for Node 2	Disk shelf 2 for Node 4	Disk shelf 2 for Node 6

Cabling the 62xx system

Correctly cabling the system enables the components to communicate with each other and with the rest of your storage environment in both normal and failure conditions.

Steps

1. [Labeling the cables](#) on page 14
You should label the cables to simplify future expansion and to make troubleshooting of the cluster easier.
2. [Cabling the disk shelf alternate control path ports](#) on page 14
The alternate control path uses the ACP ports on the disk shelves to connect to the storage controllers. Having a control path that is separate from the data path increases reliability and simplifies troubleshooting.
3. [Cabling the disk shelf SAS data ports](#) on page 15
The serial attached SCSI (SAS) ports enable each controller in the HA pair to access its own disks and the disks of its partner controller.
4. [Cabling the HA pair NVRAM ports](#) on page 22
The nonvolatile RAM (NVRAM) ports of the two controllers in an HA pair enable the pair to mirror its write cache and perform HA monitoring.
5. [Cabling the console ports](#) on page 22
The console ports enable you to manage the server when a network connection is not available, for example, during software installation or network troubleshooting.
6. [Cabling the management ports](#) on page 23
The storage controller management ports enable you to manage the cluster using the CLI or software such as System Manager. The private cluster switch management ports enable you to manage the switches.
7. [Cabling the private cluster interconnect for NetApp CN1610 switches](#) on page 24
The private cluster interconnect enables the cluster nodes to communicate with each other and to move data within the cluster. The private cluster interconnect uses dedicated 10 GbE switches.
8. [Cabling the private cluster ISLs for NetApp CN1610 switches](#) on page 26
The four inter-switch links (ISLs) between the two private cluster interconnect switches enable high availability of the cluster network.
9. [Cabling the Ethernet data networks](#) on page 26
The Ethernet ports on the storage controllers enable iSCSI initiators in your hosts to connect to the iSCSI targets in the storage controller using cables that you supply. These ports can also support other Ethernet data traffic such as NFS or CIFS.
10. [Cabling the Fibre Channel data fabrics](#) on page 27

The FC ports on the storage controllers enable FC initiators in your hosts to connect to the FC targets in the storage controller using cables that you supply. The FC ports can also be used to connect to backup devices such as tape drives.

11. *Cabling the power supplies* on page 28

The power supplies provide the electrical power needed to operate the system. The storage components should connect to two independent power sources.

Labeling the cables

You should label the cables to simplify future expansion and to make troubleshooting of the cluster easier.

Before you begin

You need the binder of labels shipped in the accessories box for the cluster.

About this task

You can label the cables at any point during the installation process.

Steps

1. Using the labels in the binder supplied with the cluster, label each end of each cable.
You do not label the power cables.
2. Save the binder containing any remaining labels for future expansion of the cluster.

Cabling the disk shelf alternate control path ports


The alternate control path uses the ACP ports on the disk shelves to connect to the storage controllers. Having a control path that is separate from the data path increases reliability and simplifies troubleshooting.

About this task

The ACP cabling uses standard CAT6 Ethernet cables.

Steps

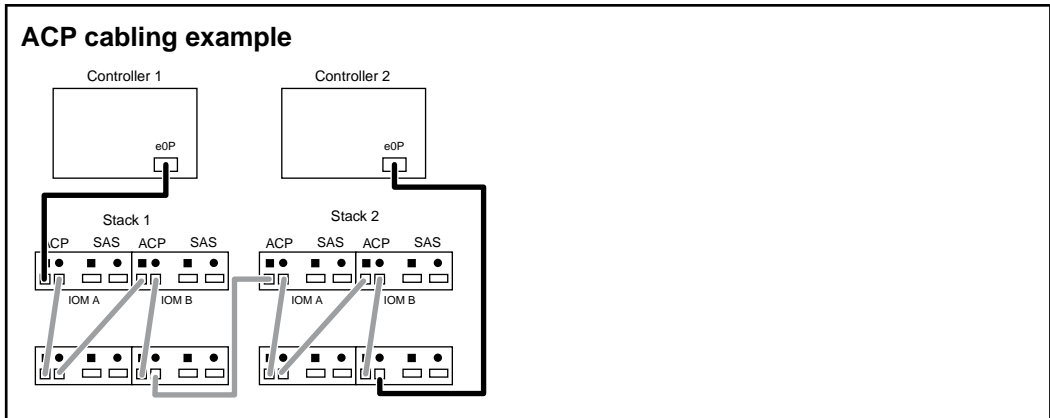
1. Connect the private management port e0P of controller 1 in an HA pair to the ACP port with the square symbol on the first disk shelf.

The e0P port is labeled with a wrench and padlock symbol ().

2. Connect the ACP port with the round symbol to the next ACP port with a square symbol.

Continue in a daisy-chain until all ACP ports for the shelves used by the HA pair are connected.

3. Connect the final ACP port with the square symbol to the private management port e0P of controller 2 in the HA pair.
4. Repeat for the other HA pairs in the cluster.



Cabling the disk shelf SAS data ports

The serial attached SCSI (SAS) ports enable each controller in the HA pair to access its own disks and the disks of its partner controller.

Before you begin

You must have located and marked which disk shelves go with each storage controller.

About this task


Disk shelf cabling is always isolated within the HA pair that owns the disks. Disk shelves are never connected to other HA pairs in the cluster.

QSFP to QSFP SAS cables are used to connect disk shelves together and to connect disk shelves to the SAS ports on the controller.

Note: It is possible to insert the QSFP connector incorrectly (upside down). The connector can slide in and even appear to click into place. However, the latch cannot engage unless the connector is inserted correctly. After inserting the cable, pull on the connector to ensure it is latched.

For disk shelf I/O Modules and controller onboard SAS ports, the cables must be inserted with the bottom of the connector facing down, as shown in the following diagram. The bottom of the

connector is slightly longer than the top, and the latch release is typically on the bottom.



For controller SAS ports on add-in cards, the orientation depends on the slot in which the SAS card is installed. Insert the connector gently, and be sure it latches securely. If the connector does not latch, rotate it 180 degrees and try again.

These instructions assume two stacks of disk shelves per HA pair. You can split your disk shelves into additional stacks to increase I/O throughput. Additional stacks require adding additional SAS ports to the storage controllers. See the *Universal SAS and ACP Cabling Guide* if you want to use additional stacks of disk shelves.

Steps

1. Set the disk shelf ID numbers for each disk shelf.

Note: The disk shelf ID numbers might have been set in the factory.

The shelf IDs in the first stack should be set to 10, 11, 12, and so on, with the first shelf ID set to 10. The second stack IDs should be 20, 21, 22, and so on. The third stack IDs should start with 30, the fourth stack IDs with 40, the fifth stack IDs with 50, and the sixth stack IDs with 60. This pattern makes it easy to keep track of each stack. IDs must be unique within an HA pair.

- a) Temporarily connect power cords to the disk shelf and turn on the power to the disk shelf.
- b) Remove the left ear cover from the shelf.
- c) Press and hold the U-shaped tab above the LEDs (DS4243) or shelf ID button (DS2246) until the first digit blinks.
- d) Press the tab or button until the correct number is displayed.
- e) Repeat steps c and d for the second digit.
- f) Press and hold the tab or button until the second number stops blinking.

Both numbers blink and the operator display panel fault LED illuminates in about five seconds.

- g) Power-cycle the disk shelf to make the new disk shelf ID take effect.
2. Starting with the first stack of disk shelves, cable the I/O modules of the stacks to each other as follows.

If your systems has only one disk shelf per stack, skip this step.

- a) Connect the IOM A SAS port with the round symbol from the first shelf to the IOM A SAS port with the square symbol on the second shelf.

IOM A is either on the top or on the left depending on your disk shelf model.


- b) Continue connecting IOM A ports on the remaining disk shelves in the stack, from the ports with a round symbol to the ports with a square symbol.

- c) Connect the IOM B SAS port with the round symbol from the first shelf to the IOM B SAS port with the square symbol on the second shelf.

IOM B is either on the bottom or on the right depending on your disk shelf model.

- d) Continue connecting IOM B ports on the remaining disk shelves in the stack, from the ports with a round symbol to the ports with a square symbol.

Example



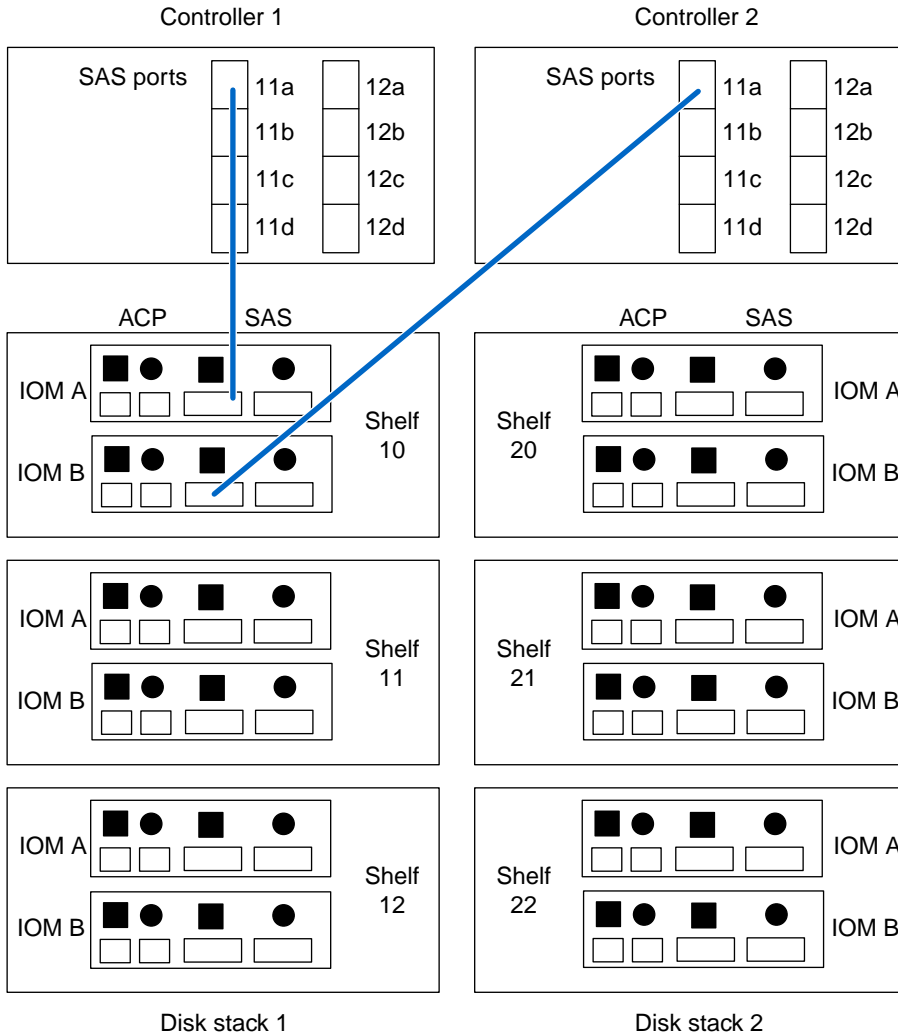
- 3. Repeat for each stack of disk shelves.

If your system has only one disk shelf per stack, skip this step.

- 4. Connect the first set of controller-to-shelf SAS cables as shown in the following diagram.

These instructions use two quad SAS adapter cards in slots 11 and 12. If your SAS cards are in different slots, adjust the port numbers for the slots actually used. For two stacks of disks, you do not use two ports on each card. However you still need two cards to avoid a single point of failure.


Example



- a) Connect controller 1 SAS port 11a to the SAS port with the square symbol in I/O module A of the first shelf in the first stack.
- b) Connect controller 2 SAS port 11a to the SAS port with the square symbol in I/O module B of the first shelf in the first stack.

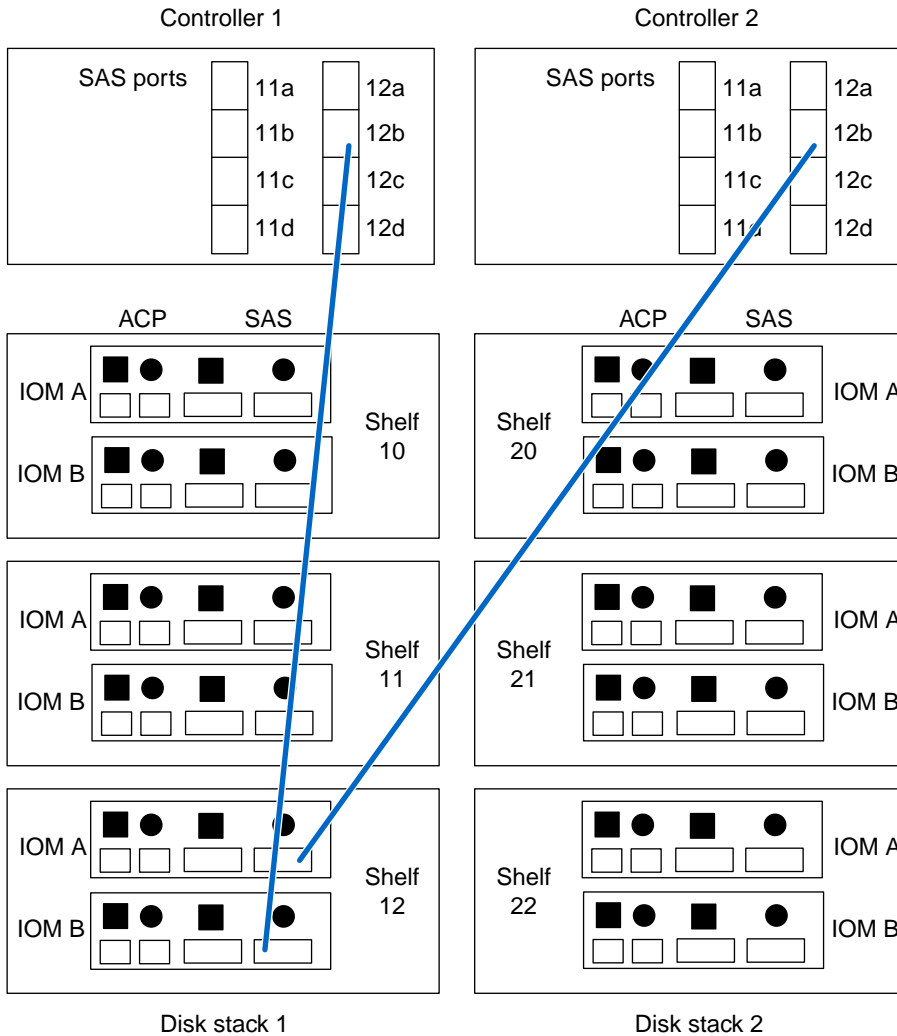
5. Connect the next set of SAS cables as shown in the following diagram.

Example




- a) Connect controller 1 SAS port 11c to the SAS port with the square symbol in I/O module A of the first shelf in the second stack.
 - b) Connect controller 2 SAS port 11c to the SAS port with the square symbol in I/O module B of the first shelf in the second stack.
6. Connect the next set of SAS cables as shown in the following diagram.

Example



- a) Connect controller 1 SAS port 12b to the SAS port with the round symbol in I/O module B of the last shelf in the first stack.
- b) Connect controller 2 SAS port 12b to the SAS port with the round symbol in I/O module A of the last shelf in the first stack.

7. Connect the final set of SAS cables as shown in the following diagram.

Example

- a) Connect controller 1 SAS port 12d to the SAS port with the round symbol in I/O module B of the last shelf in the second stack.
- b) Connect controller 2 SAS port 12d to the SAS port with the round symbol in I/O module A of the last shelf in the second stack.

8. Repeat for the other HA pairs in the cluster.

Cabling the HA pair NVRAM ports

The nonvolatile RAM (NVRAM) ports of the two controllers in an HA pair enable the pair to mirror its write cache and perform HA monitoring.


About this task

This task does not apply to systems with two controllers in the same chassis; the HA pair connection is made internally through the chassis backplane.

You use the supplied QSFP to QSFP cables or SFP modules and optical cables.

Steps

1. Connect port 2a (top port of NVRAM8 card in vertical slot 2) on the first controller in the HA pair to port 2a on the other controller.
2. Connect port 2b (bottom port of NVRAM8 card in vertical slot 2) on the first controller in the HA pair to port 2b on the other controller.
3. Repeat for the other HA pairs in the cluster.



Cabling the console ports

The console ports enable you to manage the server when a network connection is not available, for example, during software installation or network troubleshooting.

About this task

You can connect the storage controller console ports to a physical terminal or to a terminal or console server.

The console ports use an RJ-45 connector. An RJ-45 to DB-9 console adapter is included with the storage controller. You supply the appropriate cable for your terminal server or terminal.

Steps

1. Connect a cable from a terminal server or terminal to the console port of the first storage controller.

The console port is labeled with a serial data (|○|○|) symbol.

2. Repeat for each storage controller in the cluster.

After you finish

Record the mapping of terminal server ports to storage controllers.

Cabling the management ports


The storage controller management ports enable you to manage the cluster using the CLI or software such as System Manager. The private cluster switch management ports enable you to manage the switches.

About this task

You can connect the controller and cluster switch management ports to existing switches in your network or to new dedicated network switches.

Steps

1. Connect a cable from the management port of the first storage controller in an HA pair to one switch in your Ethernet network.

The management port is labeled with a wrench symbol () and is identified as e0M in commands and output.

2. Connect a cable from the management port of the other storage controller in the HA pair to a different switch in your Ethernet network.
3. Repeat for the other HA pairs in the cluster.
4. Connect a cable from the management port of the first private cluster switch to one switch in your Ethernet network.
5. Connect a cable from the management port of the second private cluster switch to a different switch in your Ethernet network.

Cabling the private cluster interconnect for NetApp CN1610 switches

The private cluster interconnect enables the cluster nodes to communicate with each other and to move data within the cluster. The private cluster interconnect uses dedicated 10 GbE switches.

About this task

The private cluster interconnect uses 10 GbE copper cables. The required cables are supplied with the system.

Additional information about the CN1610 switches is available in the *CN1601 and CN1610 Switch Setup and Configuration Guide*.

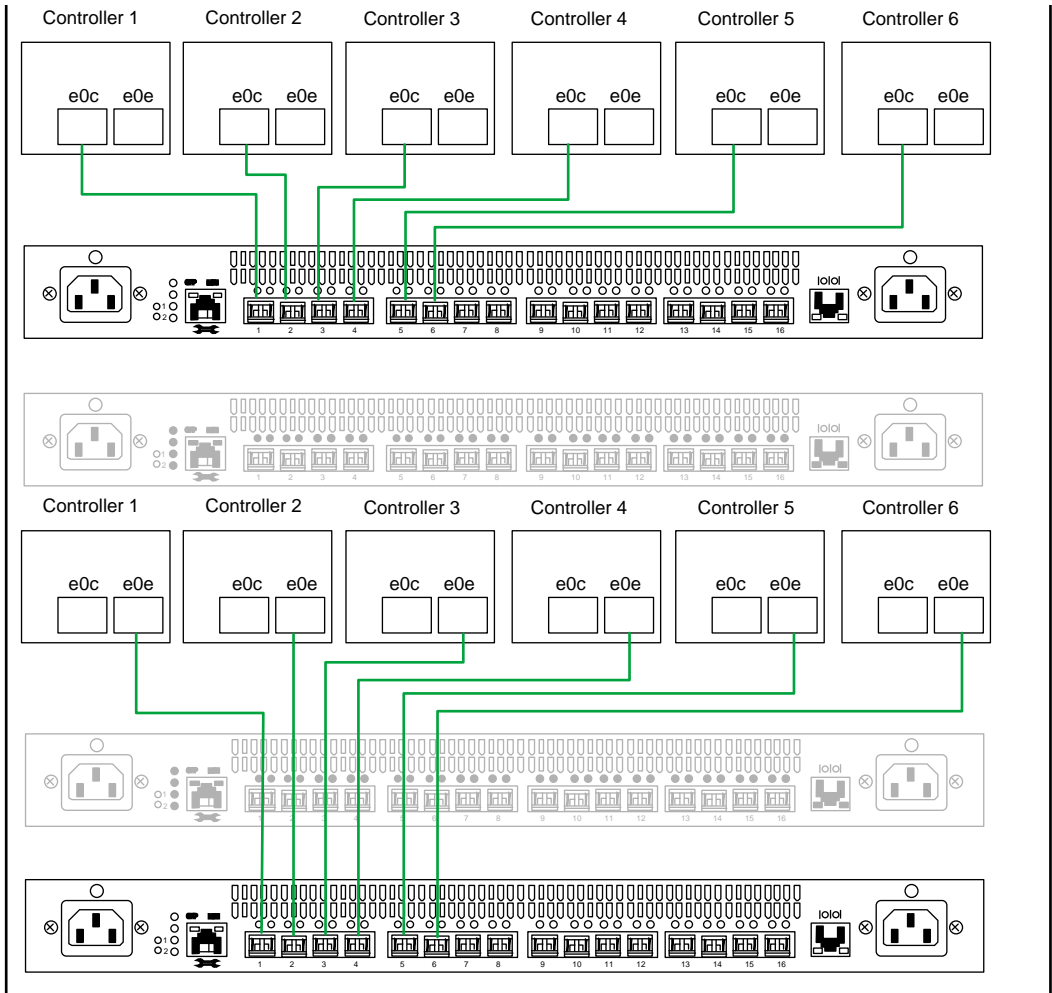
Steps

1. Using a supplied 10 GbE copper cable, connect port e0c of the first storage controller to port 1 of the top cluster switch.
2. Connect port e0e of the first storage controller to port 1 of the bottom cluster switch.
3. Repeat for the other storage controllers in the cluster.

Connect the second controller to port 2 of each switch, the third controller to port 3, and so on.

Private cluster interconnect cabling for NetApp CN1610 switches

The following example shows six nodes. You connect the cluster interconnect the same way regardless of the number of nodes in the cluster. If you are installing an eight-node cluster, controller 7 connects to port 7, and controller 8 connects to port 8.



Related information

[CN1601 and CN1610 Switch Setup and Configuration Guide at support.netapp.com/documentation/docweb/index.html?productID=61472](http://support.netapp.com/documentation/docweb/index.html?productID=61472)

Cabling the private cluster ISLs for NetApp CN1610 switches


The four inter-switch links (ISLs) between the two private cluster interconnect switches enable high availability of the cluster network.

About this task

The ISLs use 10 GbE copper cables. The required cables are supplied with the system.

Steps

1. Connect port 13 of the top cluster switch to port 13 of the bottom cluster switch.
2. Repeat for ports 14 through 16. There are 4 ISLs.



Cabling the Ethernet data networks

The Ethernet ports on the storage controllers enable iSCSI initiators in your hosts to connect to the iSCSI targets in the storage controller using cables that you supply. These ports can also support other Ethernet data traffic such as NFS or CIFS.

About this task

Connect to two or more separate networks for high availability. The specific controller ports you use depend on the controller model, bandwidth needs, and the network speed and media. Use more ports or faster ports for higher bandwidth.

Steps

1. Cable one or more Ethernet ports from the first storage controller to an Ethernet switch in your data network.

2. Cable one or more Ethernet ports from the first storage controller to a different Ethernet switch in your data network.
3. Repeat these steps for the other storage controllers in the cluster.
4. Record which storage controller ports connect to which switch ports in your network.

After you finish

Refer to the *Clustered Data ONTAP Network Management Guide* for detailed network configuration information.

Cabling the Fibre Channel data fabrics

The FC ports on the storage controllers enable FC initiators in your hosts to connect to the FC targets in the storage controller using cables that you supply. The FC ports can also be used to connect to backup devices such as tape drives.

About this task

Use two or more separate fabrics for high availability. The specific ports you use depend on the controller model.

Steps

1. Cable one or more FC ports from the first storage controller to an FC switch in your SAN.
2. Cable one or more FC ports from the first storage controller to an FC switch in a different SAN fabric.
3. Repeat these steps for the other storage controllers in the cluster.

After you finish

FC ports might default to initiator mode. For FC ports used as targets, you can change the port's mode from initiator to target after the node is running by using the `system node run fcadmin` command.

See the *Clustered Data ONTAP SAN Configuration Guide* for detailed SAN configuration information.

Cabling the power supplies

The power supplies provide the electrical power needed to operate the system. The storage components should connect to two independent power sources.

About this task

The supplied cables should fit the power receptacles in typical enclosures for your country. You might need to provide different cables.

All power supplies must be connected to a power source.

Steps

1. Connect the first power supply of each cluster component to the first power source in the enclosure or rack.
2. Connect the second power supply of each cluster component to a second power source in the enclosure or rack.

Setting up the cluster

Setting up the cluster involves creating the cluster on the first node, joining the remaining nodes to it, enabling storage failover or high availability, and synchronizing the time.

Steps

1. *Creating the cluster on the first node* on page 30
You use the Cluster Setup wizard to create the cluster on the first node. The wizard helps you to configure the cluster network that connects the nodes (if the cluster consists of two or more nodes), create the cluster admin Vserver, add feature license keys, and create the node management interface for the first node.
2. *Joining a node to the cluster* on page 31
After creating a new cluster, for each remaining node, you use the Cluster Setup wizard to join the node to the cluster and create its node management interface.
3. *Troubleshooting boot failure of a node* on page 32
If a node fails to boot when you power it on, you can check for common problems and resolve them.
4. *Configuring cluster high availability in a two-node cluster* on page 33
If the cluster contains only two nodes (a single HA pair), you should configure cluster high availability.
5. *Synchronizing the system time across the cluster* on page 34
Synchronizing the time ensures that every node in the cluster has the same time, and prevents CIFS and Kerberos failures.
6. *Reducing root volume size to prevent warning messages* on page 35
By default, each node's root volume is created using 95% of the available root aggregate space. Management software, including OnCommand System Manager, reports a warning when an aggregate is 95% full. To avoid these unnecessary warnings, you can reduce the size of the root volume.
7. *Configuring AutoSupport* on page 37
You should configure AutoSupport so that the cluster can send information to technical support to help with troubleshooting. AutoSupport messages are scanned for potential problems and are available to technical support when they assist customers in resolving issues.
8. *Configuring a node for troubleshooting and maintenance* on page 38
The service processor (SP) enables troubleshooting and maintaining a node even if the node is not running. To enable access to the SP of a node, you need to configure the SP network.

Creating the cluster on the first node

You use the Cluster Setup wizard to create the cluster on the first node. The wizard helps you to configure the cluster network that connects the nodes (if the cluster consists of two or more nodes), create the cluster admin Vserver, add feature license keys, and create the node management interface for the first node.

Before you begin

The cluster setup worksheet should be completed, the storage system hardware should be installed and cabled, and the console should be connected to the node on which you intend to create the cluster.

Steps

1. Power on the first node.

The node boots, and the Cluster Setup wizard starts on the console.

```
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.

Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

Note: If a login prompt appears instead of the Cluster Setup wizard, you must start the wizard by logging in using the factory default settings and then entering the `cluster setup` command.

2. Create a new cluster:

create

3. Follow the prompts to complete the Cluster Setup wizard:

- To accept the default value for a prompt, press Enter. The default values are determined automatically based on your platform and network configuration.
- To enter your own value for the prompt, enter the value and then press Enter.

4. After the Cluster Setup wizard is completed and exits, verify that the cluster is active and the first node is healthy:

```
cluster show
```

Example

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster1::> cluster show
Node           Health Eligibility
-----
cluster1-01    true   true
```

You can access the Cluster Setup Wizard to change any of the values you entered for the admin Vserver or node Vserver by using the `cluster setup` command.

After you finish

If the cluster consists of two or more nodes, you should join each remaining node to the cluster.

Joining a node to the cluster

After creating a new cluster, for each remaining node, you use the Cluster Setup wizard to join the node to the cluster and create its node management interface.

Before you begin

The cluster must be created on the first node.

About this task

You must complete this task for each node in the cluster.

Steps

1. Power on the node.

The node boots, and the Cluster Setup wizard starts on the console.

```
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.
```

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

2. Enter the following command to join the node to the cluster:

```
join
```

3. Follow the prompts to set up the node and join it to the cluster:

- To accept the default value for a prompt, press Enter.
- To enter your own value for the prompt, enter the value and then press Enter.

4. After the Cluster Setup wizard is completed and exits, verify that the node is healthy and eligible to participate in the cluster:

```
cluster show
```

Example

The following example shows a cluster after the second node (cluster1-02) has been joined to the cluster:

```
cluster1::> cluster show
Node           Health  Eligibility
-----
cluster1-01    true   true
cluster1-02    true   true
```

You can access the Cluster Setup Wizard to change any of the values you entered for the admin Vserver or node Vserver by using the `cluster setup` command.

5. Repeat this task for each remaining node.

Troubleshooting boot failure of a node

If a node fails to boot when you power it on, you can check for common problems and resolve them.

Before you begin

The system must be cabled and powered on.

About this task

The root aggregate for each node is loaded onto the disk shelf intended for the node. If the wrong disk shelves are connected to the node, or if the correct shelves are not correctly cabled, the node cannot boot. Each node requires one and only one root aggregate to boot.

Troubleshooting boot failures uses the maintenance mode of the node, which can be accessed using the serial console port of the node.

Steps

1. Check that all disk shelf cables appear to be connected to the correct disk shelves and that the connectors are all fully seated.
2. Verify that all disk shelves are powered on.
3. Connect a computer with a terminal emulator program to the serial console port of the node that does not boot.
4. Turn off the node that does not boot and then turn it back on, observing the startup messages on the console.
5. When the message `Press Ctrl-C for Boot Menu` appears, press Ctrl-C to display the boot menu.
6. Select option 5 at the boot menu and continue with the boot in maintenance mode.
The Maintenance mode command prompt (`*>`) is displayed.
7. Run the `aggr status` command, and then verify that exactly one aggregate is shown for the node.
8. Run the `sasadmin shelf` command, and then verify that the number of disk shelves displayed matches the number of shelves you connected to the node.
There should be two paths to each shelf, so if there are five shelves, you should see 10 entries.
9. Run the `storage show disk -p` command, and then verify that each disk is connected to two different ports.
10. If any of these tests fail, correct the disk shelf arrangement or cabling, and then rerun the command.

Configuring cluster high availability in a two-node cluster

Cluster high availability (HA) differs from the HA provided by storage failover. You must configure cluster HA in a cluster if it contains only two nodes.

About this task

Cluster HA is a special case of the clustering technology used in larger clusters; cluster HA communications occur over the cluster network, not over the HA interconnection between controllers in a storage-failover pair. In a two-node cluster, cluster HA ensures that the failure of one node does not disable the cluster.

If your cluster contains only two nodes, you must enable cluster HA. If the cluster contains more than two nodes, you should not enable cluster HA in the cluster.

Step

1. Enter the following command to enable cluster HA:

```
cluster ha modify -configured true
```

Synchronizing the system time across the cluster

Synchronizing the time ensures that every node in the cluster has the same time, and prevents CIFS and Kerberos failures.

Before you begin

A Network Time Protocol (NTP) server should be set up at your site.

About this task

You synchronize the time across the cluster by associating each node in the cluster with the NTP server. For more information about managing the system time, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Verify that the system time and time zone is set correctly for each node.

All nodes in the cluster should be set to the same time zone.

- a) Use the `cluster date show` command to display the current date, time, and time zone for each node.

Example

```
cluster1::> cluster date show
Node           Date           Timezone
-----
cluster1-01    04/06/2013 09:35:15 America/New_York
cluster1-02    04/06/2013 09:35:15 America/New_York
cluster1-03    04/06/2013 09:35:15 America/New_York
cluster1-04    04/06/2013 09:35:15 America/New_York
cluster1-05    04/06/2013 09:35:15 America/New_York
cluster1-06    04/06/2013 09:35:15 America/New_York
6 entries were displayed.
```

- b) Optional: Use the `cluster date modify` command to change the date or time zone for all of the nodes.

Example

This example changes the time zone for the cluster to be GMT:

```
cluster1::> cluster date modify -timezone GMT
```

- For each node in the cluster, use the `system services ntp server create` command to associate the node with your NTP server.

Note: The following examples assume that DNS has been configured for the cluster. If you have not configured DNS, you must specify the IP address of the NTP server.

Example

The following example associates a node named `cluster1-01` with an NTP server named `ntp1.example.com` that is running the highest-numbered version of NTP available:

```
cluster1::> system services ntp server create -node cluster1-01 -
server ntp1.example.com -version max
```

- Use the `system services ntp server show` command to verify that each node is associated with an NTP server.

Example

```
cluster1::> system services ntp server show
Node      Server                               Version
-----
cluster1-01 ntp1.example.com                    max
cluster1-02 ntp1.example.com                    max
cluster1-03 ntp1.example.com                    max
cluster1-04 ntp1.example.com                    max
cluster1-05 ntp1.example.com                    max
cluster1-06 ntp1.example.com                    max
6 entries were displayed.
```

Reducing root volume size to prevent warning messages

By default, each node's root volume is created using 95% of the available root aggregate space. Management software, including OnCommand System Manager, reports a warning when an aggregate is 95% full. To avoid these unnecessary warnings, you can reduce the size of the root volume.

About this task

Reducing the size of the root volume is not a problem because the node's root volume (`vo10`) typically has more space than needed.

This procedure must be completed using the Data ONTAP CLI.

Steps

1. Verify the size of the root volumes by using the `volume show` command.

Example

```
cluster1::> volume show -vserver * -volume vol0
-fields percent-used, size
vserver      volume size      percent-used
-----
cluster1-01 vol0    352.6GB 7%
cluster1-02 vol0    352.6GB 8%
```

2. Verify the size of the root aggregates by using the `storage aggregate show` command.

Example

```
innovate::> storage aggregate show -aggregate aggr0*
Aggregate      Size Available Used% State  #Vols  Nodes
-----
aggr0          371.9GB  17.31GB  95% online    1 cluster1-01
aggr0_cluster1_02_0
                371.9GB  17.31GB  95% online    1 cluster1-02
```

3. Compute the new root volume size as 90% of the root aggregate.

Example

For example, if the aggregate size is 371.9 GB, the volume should be reduced to 334 GB. ($371.9 * 0.9 = 334.7$)

4. Change the root volume size for each node to the new value by using the `run -node vol size` command.

Example

```
cluster1::> run -node cluster1-01 vol size vol0 334g
vol size: Flexible volume 'vol0' size set to 334g.
```

5. Verify that the root aggregates are now 90% full by using the `storage aggregate show` command.

Example

```
innovate::> storage aggregate show -aggregate aggr0*
Aggregate      Size Available Used% State  #Vols  Nodes
-----
aggr0          371.9GB  35.99GB  90% online    1 cluster1-01
```

aggr0_cluster1_02_0	371.9GB	35.99GB	90% online	1 cluster1-02
---------------------	---------	---------	------------	---------------

Configuring AutoSupport

You should configure AutoSupport so that the cluster can send information to technical support to help with troubleshooting. AutoSupport messages are scanned for potential problems and are available to technical support when they assist customers in resolving issues.

Steps

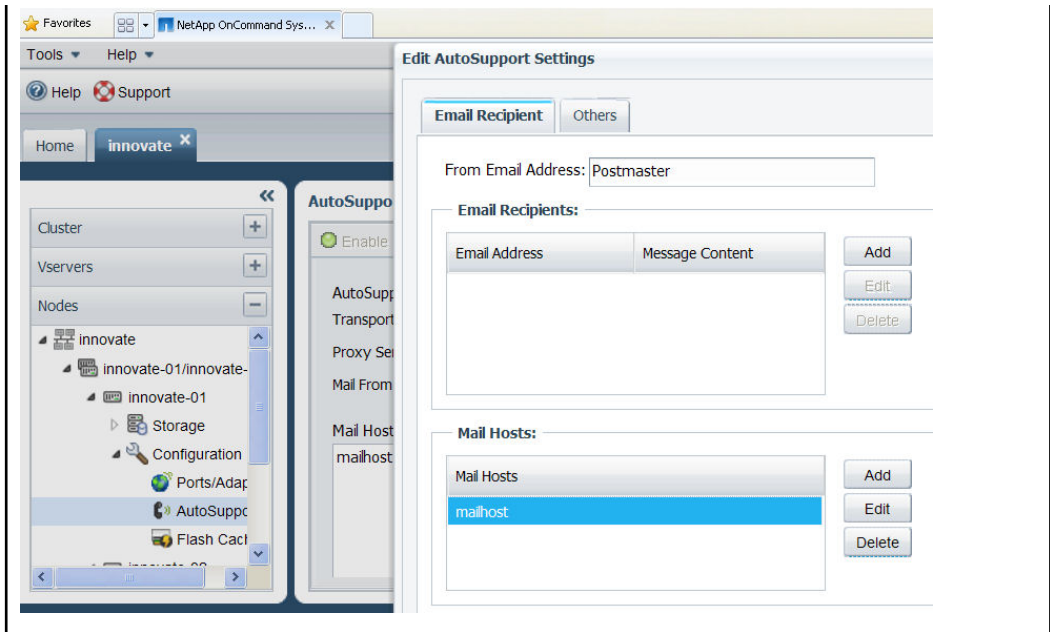
1. From the **Home** tab, double-click the appropriate cluster.
2. Expand the **Nodes** hierarchy in the left navigation pane.
3. In the navigation pane, select the node, and then click **Configuration > AutoSupport**.
4. Click **Edit**.
5. In the **E-mail Recipient** tab, type the email address from which email notifications are sent, specify the email recipients and the message content for each email recipient, and add the mail hosts.

Note: You can add up to five email addresses of the host names.

6. In the **Others** tab, select a transport protocol for delivering email messages from the drop-down list and specify the HTTP or HTTPS proxy for HTTP.
7. Click **OK**.
8. Click **Test** to send a test AutoSupport message and verify that it was received by the email recipients you specified.
9. Repeat for each node in the cluster.

AutoSupport settings

The following graphic shows the Edit AutoSupport Settings window.



Configuring a node for troubleshooting and maintenance

The service processor (SP) enables troubleshooting and maintaining a node even if the node is not running. To enable access to the SP of a node, you need to configure the SP network.

About this task

You must configure the SP network using the CLI.

For more information about configuring the SP, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Configure and enable the SP by using the `system node service-processor network modify` CLI command.
2. Display the SP network configuration to verify the settings by using the `system node service-processor network show` command.
3. Repeat for each node in the cluster.

Example of configuring the SP network

The following command configures the SP of a node:

```
cluster1::> system node service-processor network modify
-node cluster1-01 -address-type IPv4 -enable true -ip-address
192.168.123.98
-netmask 255.255.255.0 -gateway 192.168.123.1
```

Verifying cabling

You use the Config Advisor tool to verify that the system is cabled correctly.

Before you begin

The hardware setup must be complete and all components must be powered on.

About this task

If you need support for the Config Advisor tool, you must follow the procedure in the tool's online help topic “Reporting Issues in Config Advisor”. The Config Advisor tool is not supported by the typical NetApp support process.

Steps

1. Download the latest version of Config Advisor to a laptop running Windows or Linux from the NetApp Support Site.
2. Connect your laptop to the management network for the cluster.
3. If needed, change the IP address of your laptop to an unused address on the subnet for the management network.
4. Start Config Advisor and select the profile **Clustered Data ONTAP**.
5. Select the cluster switch model.
6. Enter the requested IP addresses and credentials.
7. Click **Collect Data**.

The Config Advisor tool displays any problems found. If problems are found, correct them and run the tool again.

Related information

[Config Advisor from NetApp Support Site: support.netapp.com/eservice/toolchest](http://support.netapp.com/eservice/toolchest)

Downloading and installing System Manager

You must have installed OnCommand System Manager 2.2 or later on your local computer to complete the setup and test provisioning of the cluster.

Before you begin

You must have a computer running a supported version of Windows or Linux. System requirements are listed on the NetApp Support Site where you download the System Manager software.

Steps

1. Download OnCommand System Manager 2.2 or later for your operating system from the NetApp Support Site.
2. Follow the installation instructions in the *OnCommand System Manager Quick Start Guide* to install the software.

A reboot is not required.

Related information

NetApp Support Site software download: support.netapp.com/NOW/cgi-bin/software/

Provisioning test LUNs for verifying cluster setup

Verifying that the cluster is set up correctly requires test LUNs that can be used for writing and reading data from test hosts. To create LUNs, you first must create an aggregate, a virtual storage server (Vserver), and a volume to contain the LUNs.

About this task

You can optionally delete the aggregate, Vserver, volume, and LUNs after testing is complete.

Steps

1. *Creating the test aggregate* on page 42
An aggregate organizes a set of physical disks on which you create virtual storage servers (Vservers) and volumes. A test aggregate provides the space needed for test LUNs.
2. *Creating the test Vserver* on page 43
The virtual storage server (Vserver) is the basic building block of all configurations. The Vserver contains the volume and the test LUNs and provides the FC or iSCSI target used to access the LUNs.
3. *Creating the test LUNs* on page 44
Two LUNs are needed for testing that hosts can write and read data from the cluster. The System Manager Create LUN wizard creates the necessary volumes, LUNs, and igroups, and maps the LUNs to the igroups.

Creating the test aggregate

An aggregate organizes a set of physical disks on which you create virtual storage servers (Vservers) and volumes. A test aggregate provides the space needed for test LUNs.

About this task

You can delete this aggregate after testing is complete, so the number of disks you choose is not critical. A five-disk aggregate is the fastest to create and destroy.

Note: If you have already planned the aggregate layout for the system, you can go ahead and create your first real aggregate instead of a test aggregate.

Steps

1. Start System Manager, and on the Home tab, double-click the new cluster.
The dashboard for the cluster is displayed.
2. In the navigation pane, click **Storage > Aggregates**.

3. Click **Create**.
4. On the first page of the Create Aggregate wizard, accept all of the default values.
5. On the second page of the wizard, accept all default values, including RAID size.
6. When prompted to choose the usable aggregate size, move the slider to the select 5 disks.
7. On the final wizard page, click **Create**.
8. Click **Finish**.

Creating the test Vserver

The virtual storage server (Vserver) is the basic building block of all configurations. The Vserver contains the volume and the test LUNs and provides the FC or iSCSI target used to access the LUNs.

Before you begin

You must have created an aggregate.

About this task

Do not create the Vserver or the volume in the root aggregate `aggr0`.

The Create Vserver wizard enables the SAN protocols and creates the logical interfaces (LIFs) needed for testing.

Steps

1. Start System Manager, and on the Home tab, double-click the new cluster.
The dashboard for the cluster is displayed.
2. In the navigation pane, click **Vservers**.
3. In the **Vservers** window, click **Create**.
The Create Vserver wizard opens.
4. Follow the instructions in the Create Vserver wizard.
 - a) When prompted to select protocols, select iSCSI or FC/FCoE or both.
 - b) If you plan to delete the test Vserver, you can skip the optional steps, such as configuring DNS, when prompted.
 - c) If you are using the iSCSI protocol, complete the wizard steps to configure iSCSI and create the logical interfaces.
Create two iSCSI data interfaces on two different ports per controller to avoid having a single point of failure.
 - d) If you are using the FC/FCoE protocol, complete the wizard steps to configure FC and create the logical interfaces.

Create at least two FC data interfaces on two different ports per controller to avoid having a single point of failure.

- e) After the wizard is complete, select the new Vserver in the navigation pane and view its storage, policies, and configuration.

Creating the test LUNs

Two LUNs are needed for testing that hosts can write and read data from the cluster. The System Manager Create LUN wizard creates the necessary volumes, LUNs, and igroups, and maps the LUNs to the igroups.

Before you begin

- You must have created a virtual storage server (Vserver) using System Manager.
- You must know the iSCSI node names or FC WWPNs of your test hosts.

About this task

If you do not know the iSCSI or FC identifiers of your test hosts, you can skip this step and return to it after the hosts are set up.

Steps

1. In the System Manager navigation pane, expand the Vserver you created, and then click **Storage**.
2. In the **Storage** window, click **Create LUN**.

The Create LUN wizard starts.

3. Follow the instructions in the wizard to create a volume, LUN, and igroup using the iSCSI or FC identifiers of your test host.
4. Run the LUN Create wizard again to create a second test volume and LUN for your second test host.

You can use the same igroup for both test LUNs.

Setting up test hosts and verifying system operation

By connecting test hosts and verifying you can write and read data, you can verify the storage cluster is correctly installed and ready to provision.

Before you begin

You need two supported hosts with Ethernet or FC connectivity to the storage cluster.

Steps

1. *Verifying your host configuration* on page 46
You need to verify your configuration before you connect your hosts to the storage cluster.
2. *Installing the host utilities* on page 47
The host utilities software sets timeouts and other host parameters. Installing host utilities on the test hosts ensures best performance and successful failover.
3. *Installing or enabling MPIO* on page 47
A multipath I/O (MPIO) solution is required any time a host has more than one path to a LUN. The MPIO solution can be built into the host operating system (native MPIO), or the solution can be a separate software product you install on the host.
4. *Connecting hosts to targets* on page 48
Connecting the test hosts to the storage cluster targets enables the hosts to access the test LUNs.
5. *Discovering LUNs and formatting disks* on page 48
Discovering LUNs and formatting them with a file system enables writing and reading data.
6. *Verifying setup by writing and reading data* on page 49
Verifying that you can write data to a LUN and read the data back helps ensure that the storage system is set up correctly.
7. *Testing storage failover* on page 49
If the cluster consists of more than one node, you should verify that each node can successfully fail over to another node. This helps ensure that the system is configured correctly and that you can maintain access to data if a real failure occurs.
8. *Deleting the test provisioning* on page 50
If you do not need the aggregates, virtual storage servers (Vservers), LUNs, and so on that you created to test the cluster, you can delete them. Deleting the test provisioning frees disks that you can use for your production provisioning.

Verifying your host configuration

You need to verify your configuration before you connect your hosts to the storage cluster.

Step

1. Use the Interoperability Matrix to verify that you have a supported combination of the following components:
 - Data ONTAP software version
 - Host operating system
 - Host utilities software
 - Multipath I/O (MPIO) software
 - Optional SnapDrive for Windows or SnapDrive for UNIX software
 - FC host bus adapter model, driver, and firmware
 - FC switch model and firmware version
 - iSCSI software or hardware initiator

After you search for your configuration and select a configuration name, details for that configuration are displayed in the **Configuration Details** dialog box. In the dialog box, be sure to review the information in the following tabs:

Notes	Lists important alerts and notes that are specific to your configuration. Review the alerts to identify the hotfixes that are required for your operating system.
Policies and Guidelines	Provides general guidelines for all SAN configurations. For example, you can view support information about Hyper-V in the Virtualization section, and you can view support information about third-party HBAs and CNAs in the “Third-party Host Bus Adapter (HBA) and Converged Network Adapter (CNA) Model Support” section.

Related information

[NetApp Interoperability Matrix: support.netapp.com/NOW/products/interoperability/](http://support.netapp.com/NOW/products/interoperability/)

Installing the host utilities

The host utilities software sets timeouts and other host parameters. Installing host utilities on the test hosts ensures best performance and successful failover.

Before you begin

You must have identified the supported host utilities version for your configuration by using the Interoperability Matrix.

Steps

1. Download the supported host utilities version from the NetApp Support Site.
2. Follow the instructions in the host utilities *Installation and Setup Guide* to install the software.

Related information

Host utilities software download: support.netapp.com/NOW/cgi-bin/software/

Host utilities documentation: support.netapp.com/documentation/productlibrary/index.html?productID=61343

Installing or enabling MPIO

A multipath I/O (MPIO) solution is required any time a host has more than one path to a LUN. The MPIO solution can be built into the host operating system (native MPIO), or the solution can be a separate software product you install on the host.

Before you begin

You must have identified the supported MPIO solution for your configuration by using the Interoperability Matrix.

Step

1. Follow the instructions provided with your MPIO software or host operating system to install or enable MPIO.

Connecting hosts to targets

Connecting the test hosts to the storage cluster targets enables the hosts to access the test LUNs.

About this task

You can find detailed steps in the appropriate host utilities *Installation and Setup Guide*. The exact procedure depends on your configuration.

Steps

1. Add the host initiator identifiers (WWPNs or IQNs) to the igroups of the test LUNs.

You might have already done this when you created the LUNs.

2. For Fibre Channel, create zones in the FC switches to enable the initiators to connect to the targets.

Zone by WWPN, not by physical port. Each initiator port should be in a separate zone with all of its corresponding target ports. Be sure to use the virtual (NPIV) WWPNs of the target LIFs, not the WWPNs of the physical FC ports. The LIF WWPN starts with 20 to 2f; the physical port WWPN starts with 50.

3. For iSCSI, create sessions between the host initiator and the storage target.

Discovering LUNs and formatting disks

Discovering LUNs and formatting them with a file system enables writing and reading data.

About this task

You can find detailed steps in the appropriate host utilities *Installation and Setup Guide*. The exact procedure depends on your configuration.

Steps

1. Rescan the host's disks to discover the new LUNs.
2. Format the virtual disk (LUN) with a file system.

Verifying setup by writing and reading data

Verifying that you can write data to a LUN and read the data back helps ensure that the storage system is set up correctly.

About this task

You can manually write and read files, or you can use a test tool such as the Data Test Program (dt). An advantage of using a test tool is that you can write and read data automatically while testing failover.

Steps

1. From your test hosts, write data to each LUN connected to the host.
2. Read back the data you wrote and verify that the data is correct.

Related information

Data Test Program (dt): dl.dropbox.com/u/32363629/Datatest/dt-source-v17.55.tar.gz

Testing storage failover

If the cluster consists of more than one node, you should verify that each node can successfully fail over to another node. This helps ensure that the system is configured correctly and that you can maintain access to data if a real failure occurs.

Before you begin

The cluster must consist of more than one node.

About this task

You should test storage failover on one HA pair at a time. To simplify troubleshooting if needed, do not try to fail over more than one node at a time.

For more information about storage failover, see the *Clustered Data ONTAP High-Availability Configuration Guide*.

Steps

1. Check the failover status by entering the following command:
`storage failover show`
2. Take over the node by its partner using the following command:
`storage failover takeover -ofnode nodename`

Example

```
storage failover takeover -ofnode cluster1-02
```

3. Verify that failover was completed by using the **storage failover show** command.
4. Give back the storage to the original node by using the following command:

```
storage failover giveback -ofnode nodename
```

Example

```
storage failover giveback -ofnode cluster1-02
```

5. Verify that giveback was completed by using the **storage failover show-giveback** command.
6. Revert all LIFs back to their home nodes by entering the following command:

```
network interface revert *
```

7. Repeat these steps for each remaining node in the cluster.

Deleting the test provisioning

If you do not need the aggregates, virtual storage servers (Vservers), LUNs, and so on that you created to test the cluster, you can delete them. Deleting the test provisioning frees disks that you can use for your production provisioning.

About this task

You must delete items in the order listed below. The Data ONTAP software does not allow certain items to be deleted if they still contain other items.

Steps

1. In the System Manager navigation pane, select **Vservers**, then expand the Vserver you created, and then click **Storage > LUNs**.
2. In the **LUNs** window, select each LUN, and then click **Delete**.
3. In the **LUNs** window, click the **Initiator Groups** tab, then select each initiator group, and then click **Delete**.
4. In the **LUNs** window, click the **Portsets** tab, then select each portset, and then click **Delete**.
5. In the System Manager navigation pane, click **Volumes** under the Vserver.
6. In the **Volumes** window, select each Volume, and then click **Status > Offline**.
Confirm when prompted.
7. In the **Volumes** window, select each Volume, and then click **Delete**.

Confirm when prompted.

8. In the System Manager navigation pane, select the top-level cluster entry under Vservers.
9. Select the test Vserver, and then click **Delete**.
10. In the System Manager navigation pane, select **Cluster**.
11. Expand the cluster, and then click **Storage > Aggregates**.
12. Select the test aggregate you created, and then click **Status > Offline**.

Confirm when prompted.

Note: Do not delete the root aggregates. Their names start with aggr0.

13. Select the test aggregate, and then click **Delete**.

Confirm when prompted.

The disks used for the aggregate are available as spare disks.

14. Zero all spare disks by entering the following command at the Data ONTAP command line:

```
storage disk zerospares
```

Where to go next

After testing has been completed successfully, your cluster is ready to provision. There are additional documents and tools to help you create the aggregates, virtual storage servers (Vservers), LUNs, and other elements needed for your production applications, and to complete further setup of the cluster.

Documentation references

The following documentation can help you provision your new cluster. All of the documentation is available from the NetApp Support Site.

<i>OnCommand System Manager Online Help</i>	Describes how to use System Manager to complete typical tasks. Also available as a PDF document.
<i>Clustered Data ONTAP Physical Storage Management Guide</i>	Describes how to organize disks into aggregates.
<i>Clustered Data ONTAP Logical Storage Management Guide</i>	Describes how to create volumes, LUNs, and other logical subsets of physical storage.
<i>Clustered Data ONTAP SAN Administration Guide</i>	Describes how to manage your SAN configuration.
<i>Clustered Data ONTAP SAN Configuration Guide</i>	Describes supported FC, FCoE, and iSCSI configurations and limits.
<i>Clustered Data ONTAP Network Management Guide</i>	Describes how to connect your cluster to your Ethernet networks and how to manage logical interfaces (LIFs).
<i>Clustered Data ONTAP Software Setup Guide</i>	Describes optional configuration steps, such as setting up the Event Message System (EMS) to control which system messages you receive.

Tools references

The following tool can help you provision your new cluster. The tool is available from the NetApp Support Site.

Interoperability Matrix Tool (IMT)	Lists supported combinations of hardware components, software versions, firmware, and drivers.
---	--

Related information

[Documentation on the NetApp Support Site: support.netapp.com](http://support.netapp.com)

Copyright information

Copyright © 1994–2013 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, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bypass, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- about this guide
 - deciding whether to use [5](#)
- addresses
 - required network information [6](#)
- aggregates
 - creating for test [42](#)
- AutoSupport
 - configuring [37](#)

B

- boot failures
 - troubleshooting node [32](#)

C

- cabinet [11](#)
- cables
 - labeling [14](#)
- cabling
 - console ports [22](#)
 - HA pair NVRAM ports [22](#)
 - management ports [23](#)
 - private cluster interconnect [24](#)
 - private cluster switch inter-switch links [26](#)
 - storage cluster [13](#)
 - verifying with Config Advisor [40](#)
- cluster interconnect
 - cabling [24](#)
- cluster setup
 - creating the cluster [30](#)
 - joining a node [31](#)
 - process for [29](#)
 - reducing root volume size [35](#)
 - synchronizing the system time [34](#)
- cluster switches
 - cabling inter-switch links for private cluster interconnect [26](#)
 - cabling private cluster interconnect [24](#)
- clusters
 - cabling [13](#)
 - configuring HA in two-node [33](#)
- component mounting [11](#)
- components

- considerations for unpacking and mounting [10](#)
 - Config Advisor
 - verifying cabling with [40](#)
 - configurations
 - verifying host [46](#)
 - configuring
 - AutoSupport [37](#)
 - configuring node
 - for troubleshooting and maintenance [38](#)
 - console ports
 - cabling [22](#)
 - controller
 - locating disk shelves for each [10](#)
 - Create Aggregate wizard
 - creating test aggregate using [42](#)
 - Create LUN wizard
 - for test LUNs [44](#)
 - Create Vserver wizard
 - creating test Vserver using [43](#)

D

- Data Test Program (dt)
 - for testing writing and reading [49](#)
- disk shelves
 - locating for each controller [10](#)
- disks
 - formatting for test [48](#)
- downloading
 - System Manager [41](#)

E

- enabling
 - MPIO on host [47](#)

F

- failover
 - testing storage [49](#)
- failures
 - troubleshooting node boot [32](#)

G

- getting started
 - deciding whether to use this guide [5](#)
- giveback
 - testing [49](#)

H

- HA
 - configuring in two-node clusters [33](#)
- HA pair NVRAM ports
 - cabling [22](#)
- high availability
 - configuring in two-node clusters [33](#)
- host configurations
 - verifying [46](#)
- host names
 - required network information [6](#)
- host utilities
 - installing [47](#)
- hosts
 - connecting to targets [48](#)
 - installing host utilities on test [47](#)
 - installing or enabling MPIO on [47](#)
 - prerequisite for setting up for testing [45](#)

I

- initiators
 - connecting to targets [48](#)
- installation
 - verifying your host configuration [46](#)
- installing
 - host utilities [47](#)
 - MPIO on host [47](#)
 - System Manager [41](#)
- inter-switch links
 - cabling for private cluster interconnect switches [26](#)
- interoperability matrix
 - verifying your host configuration using [46](#)
- IP addresses
 - required network information [6](#)
- ISLs
 - cabling for private cluster interconnect switches [26](#)

L

- labeling
 - cables [14](#)

LUNs

- creating for test [44](#)
- deleting test [50](#)
- discovering for test [48](#)

M

- maintenance
 - configuring the SP on a node to enable [38](#)
- maintenance mode
 - troubleshooting node boot failure [32](#)
- management ports
 - cabling [23](#)
- mounting
 - components, considerations for [10](#)
- MPIO
 - installing or enabling [47](#)
- multipath I/O
 - See* MPIO

N

- network information
 - required [6](#)
- Network Time Protocol (NTP)
 - associating nodes with [34](#)
 - enabling for the cluster [34](#)
- networks
 - configuring the SP [38](#)
- node
 - joining to a cluster [31](#)
- nodes
 - troubleshooting boot failure [32](#)
- NTP
 - See* Network Time Protocol
- NVRAM ports
 - cabling [22](#)

O

- OnCommand System Manager
 - See* System Manager
- overview
 - deciding whether to use this guide [5](#)

P

- planning
 - required network information [6](#)

- private cluster interconnect
 - cabling [24](#)
- production applications
 - where to get information about provisioning elements for [52](#)
- provisioning
 - deleting test [50](#)
 - where to get information about [52](#)

R

- rack [11](#)
- reading data
 - testing setup [49](#)
- requirements
 - network information [6](#)
 - verifying that your host configuration meets [46](#)
- root aggregates
 - troubleshooting when missing from node [32](#)
- root volume
 - identifying disk shelf containing [10](#)
- root volumes
 - reducing size [35](#)

S

- serial console ports
 - cabling [22](#)
- servers
 - prerequisite for setting up for testing [45](#)
- service processors
 - configuring the network [38](#)
- setup
 - where to get information about optional tasks [52](#)
- SPs
 - configuring the network [38](#)
- storage controller
 - locating disk shelves for each [10](#)
- switches
 - cabling inter-switch links for private cluster interconnect [26](#)
 - cabling private cluster interconnect [24](#)
- synchronizing
 - system time across the cluster [34](#)
- System Manager
 - creating test aggregate using [42](#)
 - creating test LUNs with [44](#)
 - creating test Vserver using [43](#)
 - downloading and installing [41](#)
- system time

- synchronizing [34](#)

T

- takeover
 - testing [49](#)
- targets
 - connecting host initiators [48](#)
- terminal server
 - cabling [22](#)
- test hosts
 - installing host utilities on [47](#)
- testing
 - cabling with Config Advisor [40](#)
 - connecting host initiators to targets [48](#)
 - creating aggregates for [42](#)
 - creating LUNs for [44](#)
 - creating Vserver for [43](#)
 - deleting test provisioning [50](#)
 - discovering LUNs and formatting disks [48](#)
 - installing host utilities on test host [47](#)
 - prerequisite for setting up test hosts and verifying system operation [45](#)
 - storage failover [49](#)
 - using Data Test Program (dt) [49](#)
 - writing and reading data [49](#)
- troubleshooting
 - configuring the SP on a node to enable [38](#)
 - node boot failure [32](#)

U

- unpacking
 - components, considerations for [10](#)

V

- verifying
 - writing and reading data [49](#)
- volumes
 - deleting test [50](#)
 - reducing size of root [35](#)
- Vservers
 - creating for test [43](#)
 - deleting test [50](#)

W

- WireGauge

- .See* Config Advisor
- wizards
 - creating test aggregate using Create Aggregate [42](#)
 - creating test Vserver using Create Vserver [43](#)
 - writing data
 - testing setup [49](#)