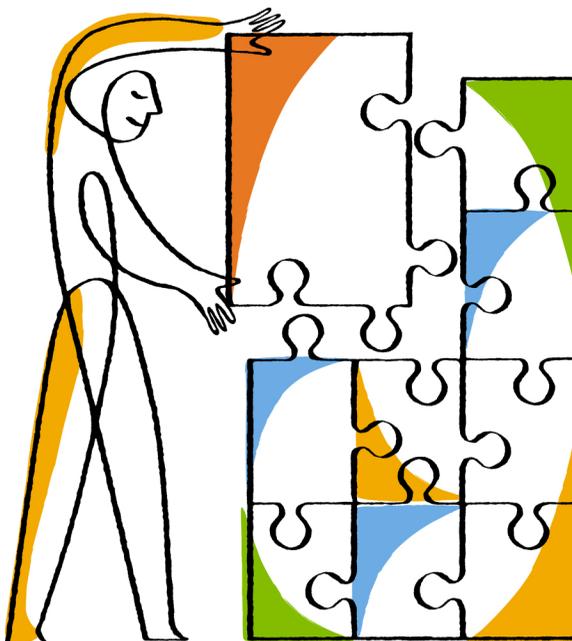




Clustered Data ONTAP® 8.2

Software Setup Guide



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: 215-07955_A0
May 2013

Contents

Preparing for the software setup process	4
Requirements for software setup	4
Registering on the NetApp Support Site	5
Completing the cluster setup worksheet	5
Completing the Vserver setup worksheet	10
Setting up the cluster	17
Creating the cluster on the first node	18
Joining a node to the cluster	19
Enabling cluster HA and switchless-cluster in a two-node cluster	20
Synchronizing the system time across the cluster	21
Setting up the configuration backup schedule for single node clusters	23
Setting up AutoSupport	24
Setting up the Event Management System	26
Setting up the Service Processor	29
Renaming a node	30
Verifying cluster setup	32
Verifying cluster health	32
Verifying that the cluster is in quorum	33
Verifying network connectivity	34
Verifying licensing	36
Verifying the high-availability configuration	37
Testing storage failover	38
Verifying the system time	39
Setting up a Vserver with FlexVol Volumes	40
Where to go from here	47
Copyright information	48
Trademark information	49
How to send your comments	50
Index	51

Preparing for the software setup process

Before setting up the software, you must complete the software setup prerequisites, register on the NetApp Support Site, and gather cluster and virtual storage server (Vserver) configuration information.

Requirements for software setup

Before you begin the software setup process, you should ensure that you have met the training, site, and installation requirements.

Training requirements

You should have completed the NetApp training program specific to your role. For more information, see the Customer Learning Map at the NetApp LearningCenter.

Site requirements

Your site must meet the physical, connectivity, power, and model-specific requirements for your cluster. For more information, see the *Site Requirements Guide*.

Hardware installation requirements

Each of the following hardware components should be installed:

- The controllers and disk shelves should be racked and cabled according to the *Installation and Setup Instructions* for your platform and the *Clustered Data ONTAP High-Availability Configuration Guide*.
- If you are setting up a switched cluster, the cluster management and interconnect switches should be installed and configured according to the *Clustered Data ONTAP Switch Setup Guide for Cisco Switches*.
- If you are using third-party storage, then you should have reviewed the *V-Series Installation Requirements and Reference Guide*, and the *V-Series Implementation Guide for Third-Party Storage*.
- The serial console should be connected to the cluster.

Related information

The NetApp Support Site: support.netapp.com

Customer Learning Map: learningcenter.netapp.com/content/public/production/learning_maps/customer/lm_customer_t1.html

Registering on the NetApp Support Site

Registering on the NetApp Support Site involves creating a new user account and registering your installed products. After registering, you can access customized support information for your system, find troubleshooting information and product documentation, download software and firmware, and request technical assistance.

About this task

You should plan for one business day for your new user account request to be processed. For additional information and best practices about the NetApp Support Site, see the *NetApp Support Owner's Manual*.

Steps

1. Go to the NetApp Support Site at support.netapp.com.
2. Click **Register Now**, and follow the instructions on the page to register as a new user.
You will be notified by email when your registration request has been processed. This process takes about a day.
3. Click **My Support > Register Products**, and follow the instructions on the page to register your new system.
Registering your system ensures that NetApp can provide you with support for your installed products.

Related information

NetApp Support Owner's Manual: support.netapp.com/NOW/products/globalservices

Completing the cluster setup worksheet

Use this worksheet to record the values that you need during the cluster setup process. If a default value is provided, you can use that value or else enter your own.

System defaults (for clusters with 2 or more nodes)

The system defaults are the default values for the private cluster network. It is best to use these default values. However, if they do not meet your requirements, you can use the table below to record your own values.

You only need to consider the system defaults for clusters with two or more nodes. Single node clusters do not use a cluster network.

Types of information	Default	Your values
Private cluster network ports	See the <i>Clustered Data ONTAP Network Management Guide</i> .	
Cluster ports' MTU size Every node in the cluster must have the same MTU size as the cluster interconnect switches.	9000 bytes	
Cluster network netmask	255.255.0.0	
Cluster interface IP addresses (for each cluster network port on each node) The IP addresses for each node must be on the same subnet.	169.254.x.x	
Switched configuration If you are setting up a two-node cluster, you can set up the cluster in either a two-node switchless configuration without cluster interconnect switches, or a switched configuration that uses the cluster interconnect switches.	Set at the factory	

Cluster information

Types of information	Your values
Cluster name The name must begin with a letter, and it must be fewer than 44 characters. The name can include the following special characters: ".", "-", and "_".	
Cluster base license key To get this license key, go to the NetApp Support Site at support.netapp.com and click My Support > Software Licenses .	

Feature license keys

You can find license keys for your initial or add-on software orders at the NetApp Support Site under **My Support > Software Licenses**. For instance, you can search with the serial number of a node to find all license keys associated with the node. Your search results will include license information for all nodes in the cluster. You can also search by cluster serial number or sales order number. If you cannot locate your license keys from the Software Licenses page, you should contact your sales or support representative.

Types of information	Your values
Feature license keys	

Admin Vserver

Types of information	Your values
<p>Cluster administrator password</p> <p>The password for the admin account that the storage system requires before granting cluster administrator access to the console or through a secure protocol.</p> <p>The default rules for passwords are as follows:</p> <ul style="list-style-type: none"> • A password must be at least eight characters long. • A password must contain at least one letter and one number. 	
<p>Cluster management interface port</p> <p>The physical port that is connected to the data network and enables the cluster administrator to manage the cluster.</p> <p>Because the cluster management interface can fail over to any node in the cluster, the cluster management interface port should have a port role of data.</p>	
<p>Cluster management interface IP address</p> <p>A unique IP address for the cluster management interface. The cluster administrator uses this address to access the admin Vserver and manage the cluster. Typically, this address should be on the data network.</p> <p>You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.</p> <p>Example: 192.0.2.66</p>	
<p>Cluster management interface netmask</p> <p>The subnet mask that defines the range of valid IP addresses on the cluster management network.</p> <p>Example: 255.255.255.0</p>	
<p>Cluster management interface default gateway</p> <p>The IP address for the router on the cluster management network.</p>	

Types of information	Your values
<p>DNS domain name</p> <p>The name of your network's DNS domain.</p> <p>The domain name must consist of alphanumeric characters. To enter multiple DNS domain names, separate each name with either a comma or a space.</p>	
<p>Name server IP addresses</p> <p>The IP addresses of the DNS name servers.</p> <p>Separate each address with either a comma or a space.</p>	

Node information (for each node in the cluster)

Types of information	Your values
<p>Physical location of the controller</p> <p>A description of the physical location of the controller. Use a description that identifies where to find this node in the cluster (for example, "Lab 5, Row 7, Rack B").</p>	
<p>Node management interface port</p> <p>The physical port that is connected to the node management network and enables the cluster administrator to manage the node.</p> <p>Because the node management interface does not fail over, the node management interface port should typically have a port role of node management; however, if necessary, it can reside on a data port.</p>	
<p>Node management interface IP address</p> <p>A unique IP address for the node management interface on the management network. If you defined the node management interface port to be a data port, then this IP address should be a unique IP address on the data network.</p> <p>You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.</p> <p>Example: 192.0.2.66</p>	

Types of information	Your values
<p>Node management interface netmask</p> <p>The subnet mask that defines the range of valid IP addresses on the node management network. If you defined the node management interface port to be a data port, then the netmask should be the subnet mask for the data network.</p> <p>Example: 255.255.255.0</p>	
<p>Node management interface default gateway</p> <p>The IP address for the router on the node management network.</p>	

NTP server information

Types of information	Your values
<p>NTP server address</p> <p>The IP address(es) of the Network Time Protocol (NTP) server at your site. This server is used to synchronize the time across the cluster.</p>	

Completing the Vserver setup worksheet

Before you start the Vserver Setup wizard to create and configure a Vserver, you must gather the required information to complete the wizard successfully.

Note: You can create and configure only Vservers with FlexVol volumes by using the Vserver Setup wizard.

The Vserver Setup wizard has the following subwizards, which you can run after you create a Vserver:

- Network setup
- Storage setup
- Services setup
- Data access protocol setup

Each subwizard has its specific requirements, depending on the types of services, protocols, and the protocol traffic.

You can use the following worksheet to record values for the setup process:

Vserver information

Types of information	Your values
<p><i>Vserver name</i></p> <p>The name of a Vserver can contain alphanumeric characters and the following special characters: ".", "-", and "_". However, the name of a Vserver should not start with a number or the following special characters: "." and "-".</p> <p>The maximum number of characters allowed in a Vserver name is 47.</p> <p>Note: Vserver names must be unique. You must use the fully qualified domain name (FQDN) of the Vserver or another convention that ensures unique Vserver names.</p>	
<p><i>Data protocols</i></p> <p>Protocols that you want to configure or allow on that Vserver</p>	
<p><i>Client services</i></p> <p>Services that you want to configure on the Vserver</p>	
<p><i>Aggregate name</i></p> <p>Aggregate on which you want to create the Vserver's root volume. The default aggregate name is used if you do not specify one.</p>	
<p><i>Language setting</i></p> <p>The default language 'C.UTF-8' is used if you do not specify one.</p> <p>The language is set for a Vserver. The language of the Vserver determines default language setting for volumes in that Vserver.</p> <p>Note: The language of a Vserver is inherited by its volumes if the language is not specified when creating the volumes.</p> <p>For more information about the available language options, see the <i>Clustered Data ONTAP System Administration Guide for Cluster Administrators</i>.</p>	

Types of information	Your values
<p><i>Vserver root volume's security style</i></p> <p>Determines the type of permissions that can be used to control data access to a volume</p> <p>For more information about the security styles, see the <i>Clustered Data ONTAP File Access and Protocols Management Guide</i>.</p>	

Information for creating volumes on the Vserver

Types of information	Values
<p><i>Volume name</i></p> <p>The default volume name is used if you do not specify one.</p>	
<p><i>Aggregate name</i></p> <p>Aggregate on which you want to create the volume. The default aggregate name is used if you do not specify one.</p>	
<p><i>Volume size</i></p>	
<p><i>Volume junction path</i></p> <p>The default junction path is used if you do not specify one.</p>	

Information for creating an IP network interface on the Vserver

Types of information	Values
<p><i>LIF name</i></p> <p>The default LIF name is used if you do not specify one.</p>	
<p><i>Protocols</i></p> <p>Protocols that can use the LIF</p> <p>Note: Protocols that can use the LIF cannot be modified after the LIF is created.</p>	
<p><i>Home node</i></p> <p>Home node is the node on which you want to create a LIF. The default home node is used if you do not specify one.</p>	

Types of information	Values
<p><i>Home port</i></p> <p>Home port is the port on which you want to create a LIF. The default home port is used if you do not specify one.</p>	
<p><i>IP address</i></p>	
<p><i>Network mask</i></p>	
<p><i>Default gateway IP address</i></p>	

Information for creating an FC network interface on the Vserver

Types of information	Values
<p><i>LIF name</i></p> <p>The default LIF name is used if you do not specify one.</p>	
<p><i>Protocols</i></p> <p>Protocols that can use the LIF</p> <p>Note: Protocols that can use the LIF cannot be modified after the LIF is created.</p>	
<p><i>Home node</i></p> <p>Home node is the node on which you want to create a LIF. The default home node is used if you do not specify one.</p>	
<p><i>Home port</i></p> <p>Home port is the port on which you want to create a LIF. The default home port is used if you do not specify one.</p>	

Information for configuring LDAP

Types of information	Values
<p><i>LDAP server IP address</i></p>	
<p><i>LDAP server port number</i></p> <p>The default LDAP server port number is used if you do not specify one.</p>	
<p><i>LDAP server minimum bind authentication level</i></p>	

Types of information	Values
<i>Bind domain name and password</i>	
<i>Base domain name</i>	

Information for configuring NIS

Types of information	Values
<i>NIS domain name</i>	
<i>IP addresses of the NIS servers</i>	

Information for configuring DNS

Types of information	Values
<i>DNS domain name</i>	
<i>IP addresses of the DNS servers</i>	

Note: You do not need to enter any information to configure NFS on a Vserver. The NFS configuration is created when you specify the protocol value as `nfs`.

Information for configuring CIFS protocol

Types of information	Values
<i>Domain name</i>	
<i>CIFS share name</i> The default CIFS share name is used if you do not specify one. Note: You must not use space characters or Unicode characters in CIFS share names. You can use alphanumeric characters and any of the following special characters: ! @ # \$ % & () _ ' { } . ~ -.	
<i>CIFS share path</i> The default CIFS share path is used if you do not specify one.	

Types of information	Values
<p><i>CIFS access control list</i></p> <p>The default CIFS access control list is used if you do not specify one.</p>	

Information for configuring iSCSI protocol

Types of information	Values
<p><i>igroup name</i></p> <p>The default igroup name is used if you do not specify one.</p>	
<p><i>Names of the initiators</i></p>	
<p><i>Operating system type of the initiator</i></p>	
<p><i>LUN name</i></p> <p>The default LUN name is used if you do not specify one.</p>	
<p><i>Volume for LUN</i></p> <p>Volume that is to be used for the LUN</p>	
<p><i>LUN size</i></p>	

Information for configuring Fibre Channel (FC) protocol (FCoE included)

Types of information	Values
<p><i>igroup name</i></p> <p>The default igroup name is used if you do not specify one.</p>	
<p><i>World wide port number (WWPN) of the initiators</i></p>	
<p><i>Operating system type of the initiator</i></p>	
<p><i>LUN name</i></p> <p>The default LUN name is used if you do not specify one.</p>	
<p><i>Volume for LUN</i></p> <p>Volume that is to be used for the LUN</p>	
<p><i>LUN size</i></p>	

Setting up the cluster

Setting up the cluster involves creating the cluster on the first node, joining any remaining nodes to the cluster, and configuring a number of features—such as synchronizing the system time—that enable the cluster to operate nondisruptively.

Steps

1. *Creating the cluster on the first node* on page 18
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 19
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. *Enabling cluster HA and switchless-cluster in a two-node cluster* on page 20
If the cluster contains only two nodes (a single HA pair), you should configure cluster high availability.
4. *Synchronizing the system time across the cluster* on page 21
Synchronizing the time ensures that every node in the cluster has the same time, and prevents CIFS and Kerberos failures.
5. *Setting up the configuration backup schedule for single node clusters* on page 23
If the cluster consists of a single node, you should set up the configuration backup schedule to back up the cluster configuration to a remote URL. This ensures that you can recover the cluster's configuration even if the node is inaccessible.
6. *Setting up AutoSupport* on page 24
You can control whether and how AutoSupport information is sent to NetApp technical support and your internal support organization, and then test that the configuration is correct.
7. *Setting up the Event Management System* on page 26
You can configure EMS to reduce the number of event messages that you receive, and to set up the event destinations and the event routes for a particular event severity.
8. *Setting up the Service Processor* on page 29
Before you can access the SP of a node, the SP network must be configured and enabled. You can configure the SP to use IPv4, IPv6, or both. The SP IPv4 configuration supports static and DHCP addressing, and the SP IPv6 configuration supports static addressing only.
9. *Renaming a node* on page 30
You can change a node's name as needed.

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.

Enabling cluster HA and switchless-cluster in a two-node cluster

A cluster consisting of only two nodes requires special configuration settings. Cluster high availability (HA) must be configured in a cluster if it contains only two nodes and it differs from the HA provided by storage failover. Also, if you have a switchless configuration, the `switchless-cluster` option must be enabled.

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.

If you have a two-node switchless configuration in which there is no cluster interconnect switch, you must ensure that the `switchless-cluster-network` option is enabled. This ensures proper cluster communication between the nodes.

Steps

1. Enter the following command to enable cluster HA:

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

2. If you have a two-node switchless cluster, enter the following commands to verify that the `switchless-cluster` option is set:

- a) Enter the following command to change to the advanced-privilege level:

```
set -privilege advanced
```

You can respond "y" when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

- b) Enter the following command:

```
network options switchless-cluster show
```

If the output shows that the value is false, you must issue the following command:

```
network options switchless-cluster modify true
```

- c) Enter the following command to return to the admin privilege level:

```
set -privilege admin
```

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
```

2. 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
```

3. 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.
```

Setting up the configuration backup schedule for single node clusters

If the cluster consists of a single node, you should set up the configuration backup schedule to back up the cluster configuration to a remote URL. This ensures that you can recover the cluster's configuration even if the node is inaccessible.

Before you begin

You must have HTTP or FTP access to a web server with PUT operations enabled.

About this task

For more information about backing up and restoring the cluster configuration, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. Set the privilege level to advanced:
`set -privilege advanced`
2. Use the `system configuration backup settings modify` command to set up the configuration backup schedule.

You can specify a remote URL where the cluster configuration backups will be uploaded, a user name for logging in to the URL, and the number of backups to keep for each configuration backup schedule. For more information, see the man page.

Example

This example specifies a remote URL to which the cluster configuration will be backed up, and the user name required to log in to the URL and upload the configuration backup file.

```
cluster1::*> system configuration backup settings modify -
destination ftp://www.example.com/config/uploads/ -username admin
```

3. If the user name that you specified for the URL requires a password, use the `system configuration backup settings set-password` command to specify the password.
4. To verify that the cluster configuration backup files can be uploaded to the remote URL, create a test configuration backup file:

```
system configuration backup create -node node_name -backup-name configuration_backup_name -backup-type cluster
```

Example

```
cluster1::*> system configuration backup create -node cluster1-01 -
backup-name test_config_backup -backup-type cluster
[Job 3592] Job is queued: Cluster Backup OnDemand Job.
```

5. Verify that the test configuration backup file can be uploaded to the remote URL:

```
system configuration backup upload -node node_name -backup configuration_backup_name -destination remote_URL
```

Example

```
cluster1::*> system configuration backup upload -node cluster1-01 -
backup test_config_backup.7z -destination ftp://www.example.com/
config/uploads/testconfig
```

6. Return to the admin privilege level:

```
set -privilege admin
```

Setting up AutoSupport

You can control whether and how AutoSupport information is sent to NetApp technical support and your internal support organization, and then test that the configuration is correct.

About this task

Perform this procedure on each node in your system where you want to configure AutoSupport.

For more information about the following commands, see the man pages.

Steps

1. Ensure AutoSupport is enabled by setting the `-state` parameter of the `system node autosupport modify` command to enable.
2. If you want technical support to receive AutoSupport messages, set the following parameters of the `system node autosupport modify` command:
 - a) Set `-support` to enable.

- b) Select a transport protocol for messages to NetApp technical support by setting `-transport` to `smtp`, `http`, or `https`.
 - c) If you chose HTTP or HTTPS as the transport protocol and you use a proxy, set `-proxy-url` to the URL of your proxy.
3. If you want your internal support organization or a support partner to receive AutoSupport messages, perform the following actions:
- a) Identify the recipients in your organization by setting the following parameters of the `system node autosupport modify` command:

Set this parameter	To this
<code>-to</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive key AutoSupport messages
<code>-noteto</code>	Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive a shortened version of key AutoSupport messages designed for cell phones and other mobile devices
<code>-partner-address</code>	Up to five comma-separated individual email addresses or distribution lists in your support partner organization that will receive all AutoSupport messages

- b) Check that addresses are correctly configured by listing the destinations using the `system node autosupport destinations show` command.
4. If you are sending messages to your internal support organization or you chose SMTP transport for messages to technical support, configure SMTP by setting the following parameters of the `system node autosupport modify` command:
- Set `-mail-hosts` to one or more mail hosts, separated by commas. You can set a maximum of five.
 - Set `-from` to the email address that sends the AutoSupport message.
 - Set `-max-smtp-size` to the email size limit of your SMTP server.
5. If you want AutoSupport to specify a fully qualified domain name when it sends connection requests to your SMTP mail server, configure DNS.

For information about configuring DNS, see the *Clustered Data ONTAP Network Management Guide*.

6. Optional: Change the following settings:

If you want to do this...	Set the following parameters of the <code>system node autosupport modify</code> command...
Hide private data by removing, masking, or encoding sensitive data in the messages	Set <code>-remove-private-data</code> to <code>true</code> . Note: If you change from <code>false</code> to <code>true</code> , all AutoSupport history and all associated files are deleted.
Stop sending performance data in periodic AutoSupport messages	Set <code>-perf</code> to <code>false</code> .

7. Check the overall configuration using the `system node autosupport show` command with the `-node` parameter.
8. Test that AutoSupport messages are being sent and received:
 - a) Use the `system node autosupport invoke` command with the `-type` parameter set to `test`.

Example

```
cluster1::> system node autosupport invoke -type test -node node1
```

- b) Confirm that NetApp is receiving your AutoSupport messages by checking the email address that technical support has on file for the system owner, who should have received an automated response from the NetApp mail handler.
- c) Optional: Confirm that the AutoSupport message is being sent to your internal support organization or to your support partner by checking the email of any address that you configured for the `-to`, `-noteto`, or `-partner-address` parameters of the `system node autosupport modify` command.

Setting up the Event Management System

You can configure EMS to reduce the number of event messages that you receive, and to set up the event destinations and the event routes for a particular event severity.

Steps

1. To see what is currently configured for the mail locations, enter the following command:
`event config show`

Example

The following command shows the configured mail locations:

```
cluster1::> event config show

Mail From: admin@localhost
Mail Server: localhost
```

2. If you need to change the mail locations, enter the following command:

```
event config modify -mailserver name -mailfrom email address
```

Example

The following example shows how to change the mail locations and display the results:

```
cluster1::> event config modify -mailserver mailhost.example.com
-mailfrom admin@node1-example.com

cluster1::> event config show

Mail From: admin@node1-example.com
Mail Server: mailhost.example.com
```

3. To create the destination for events, enter the following command and specify the name and email address:

```
event destination create -name destination -mail email address
```

You can send events to email addresses, SNMP trap hosts, and syslog servers.

Example

The following command creates an email destination and sends all important events to the specified email address:

```
cluster1::> event destination create -name test_dest -mail
me@example.com
```

4. Use the `event route add-destinations` command to define the severity level of messages to receive.

The recommended practice is to set up event routes for critical and above events.

Example

The following example sends all critical, alert, and emergency events to the `test_dest` event destination, and displays the results:

```
cluster1::> event route add-destinations {-severity <=CRITICAL}
```

```
-destinations test_dest
```

```
cluster1::> event dest show
```

Name	Mail Dest.	SNMP Dest.	Syslog Dest.	Hide Params
allevents	-	-	-	false
asup	-	-	-	false
criticals	-	-	-	false
pager	-	-	-	false
test_dest	me@example.com	-	-	false
traphost	-	-	-	false

- To display all critical and above events, enter the following command:

```
event route show -severity <=CRITICAL
```

Example

The following example shows the events with critical and above severity levels:

```
cluster1::> event route show -severity -CRITICAL
```

Message Threshd	Severity	Destinations	Freq Threshd	Time
adminapi.time.zoneDiff	ALERT	test_dest	0	3600
api.engine.killed	CRITICAL	test_dest	0	0
app.log.alert	ALERT	test_dest	0	0
app.log.crit	CRITICAL	test_dest	0	0
app.log.emerg	EMERGENCY	test_dest	0	0

- If you are still getting too many event messages, use the `-timethreshold` option to specify how often events are sent to the destination.

Example

For example, the following event is displayed once per hour:

```
cluster1::> event route modify -messagename adminapi.time.zoneDiff
-timethreshold 3600
```

Result

When you have completed these steps, all critical events are automatically sent to the destination specified in the event route.

Setting up the Service Processor

Before you can access the SP of a node, the SP network must be configured and enabled. You can configure the SP to use IPv4, IPv6, or both. The SP IPv4 configuration supports static and DHCP addressing, and the SP IPv6 configuration supports static addressing only.

Before you begin

To configure IPv6 connections for the SP, IPv6 must already be configured and enabled for Data ONTAP. The `network options ipv6` commands manage IPv6 settings for Data ONTAP. For more information about IPv6 configuration, see the *Clustered Data ONTAP Network Management Guide*.

Steps

1. Configure and enable the SP by using the `system node service-processor network modify` command.
 - The `-address-type` parameter specifies whether the IPv4 or IPv6 configuration of the SP is to be modified.
 - The `-enable` parameter enables the network interface of the specified IP address type.
 - The `-dhcp` parameter specifies whether to use the network configuration from the DHCP server or the network address that you provide.
You can enable DHCP (by setting `-dhcp` to `v4`) only if you are using IPv4. You cannot enable DHCP for IPv6 configurations.
 - The `-ip-address` parameter specifies the public IP address for the SP.
 - The `-netmask` parameter specifies the netmask for the SP (if using IPv4.)
 - The `-prefix-length` parameter specifies the network prefix-length of the subnet mask for the SP (if using IPv6.)
 - The `-gateway` specifies the gateway IP address for the SP.

For more information about the `system node service-processor network modify` command, see the man page.

2. Display the SP network configuration to verify the settings by using the `system node service-processor network show` command.

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

Example of configuring the SP network

The following example configures the SP of a node to use IPv4, enables the SP, and displays the SP network configuration to verify the settings.

```

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

cluster1::> system node service-processor network show -instance -node local

                Node: node1
                Address Type: IPv4
                Interface Enabled: true
                Type of Device: SP
                  Status: online
                Link Status: up
                DHCP Status: none
                IP Address: 192.168.123.98
                MAC Address: ab:cd:ef:fe:ed:02
                Netmask: 255.255.255.0
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: 192.168.123.1

                Node: node1
                Address Type: IPv6
                Interface Enabled: false
                Type of Device: SP
                  Status: online
                Link Status: disabled
                DHCP Status: none
                IP Address: -
                MAC Address: ab:cd:ef:fe:ed:02
                Netmask: -
Prefix Length of Subnet Mask: -
Router Assigned IP Address: -
Link Local IP Address: -
Gateway IP Address: -
2 entries were displayed.

cluster1::>

```

Renaming a node

You can change a node's name as needed.

Step

1. To rename a node, use the `system node rename` command.

The maximum length of a node's name is 47 characters.

Example

The following command renames node “node1” to “node1a”:

```
cluster1::> system node rename -node node1 -newname node1a
```

Verifying cluster setup

Misconfiguring the cluster during cluster setup can result in errors that are difficult to troubleshoot. Accordingly, after setting up the cluster, you should complete verification tasks to ensure that the cluster is operational and configured according to your requirements.

Verifying cluster health

After completing cluster setup, you should verify that each node is healthy and eligible to participate in the cluster.

About this task

For more information about node health and eligibility, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Step

1. Use the `cluster show` command to view the status of each node.

Example

This example shows that each node is healthy and eligible as indicated by status `true` in the `Health` and `Eligibility` columns (Status `false` indicates a problem).

```
cluster1::> cluster show
Node           Health  Eligibility
-----
node0          true   true
node1          true   true
node2          true   true
node3          true   true
4 entries were displayed.
```

Verifying that the cluster is in quorum

After setting up the cluster, you must ensure that all nodes are participating in a replicated database (RDB) quorum and that all rings are in the quorum. You must also verify that the per-ring quorum master is the same for all nodes.

About this task

For more information about cluster replication rings and RDB quorums, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. At the advanced privilege level, display each RDB process:

To display this RDB process...	Enter this command...
Management application	<code>cluster ring show -unitname mgmt</code>
Volume location database	<code>cluster ring show -unitname vlodb</code>
Virtual-Interface manager	<code>cluster ring show -unitname vifmgr</code>
SAN management daemon	<code>cluster ring show -unitname bcomd</code>

Example

```
cluster1::*> cluster ring show -unitname vlodb
Node   UnitName Epoch DB Epoch DB Trnxs Master
-----
node0  vlodb    154   154   14847  node0
node1  vlodb    154   154   14847  node0
node2  vlodb    154   154   14847  node0
node3  vlodb    154   154   14847  node0
4 entries were displayed.
```

For each process, verify the following configuration details:

- The relational database epoch and database epochs match for each node.
- The per-ring quorum master is the same for all nodes.
Note that each ring might have a different quorum master.

2. If you are operating in a SAN environment, verify that the cluster is in a SAN quorum:

```
event log show -messagename scsiblade.*
```

The `scsiblade` event message should indicate that the `scsi-blade` is in quorum.

Example

```
cluster::> event log show -messagename scsiblade.*
Time           Node           Severity      Event
-----
8/13/2012 14:03:51 node0          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
8/13/2012 14:03:51 node1          INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...
```

Verifying network connectivity

You should verify that the cluster, cluster management, and node management interfaces are configured correctly.

Steps

1. If the cluster has more than one node, at the advanced privilege level, use the `cluster ping-cluster` command to ping all combinations of the cluster LIFs from each node.

If the cluster consists of a single node, you should skip this step.

Example

This example pings the cluster LIFs from node1.

```
cluster1::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.254.231.102  10.254.91.42
Remote = 10.254.42.25   10.254.16.228
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 4 path(s):
  Local 10.254.231.102 to Remote 10.254.16.228
  Local 10.254.231.102 to Remote 10.254.42.25
  Local 10.254.91.42 to Remote 10.254.16.228
  Local 10.254.91.42 to Remote 10.254.42.25
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

Complete this step for each node in the cluster. For each node, you should verify the following:

- All of the paths are up.
- The pings are successful at each MTU size (1500, 4500, and 9000).
If the pings are only successful for MTU size 1500, then verify that the cluster network switch and cluster ports are configured with the correct MTU sizes. For more information about

configuring the MTU size for a port, see the *Clustered Data ONTAP Network Management Guide*.

2. Use the `network interface show` command to verify that the cluster management and node management LIFs are configured correctly.

Example

```
cluster1::> network interface show
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
cluster1	cluster_mgmt	up/up	172.17.178.119/24	ie3070-1	e1a	true
node0	clus1	up/up	172.17.177.120/24	ie3070-1	e0a	true
	clus2	up/up	172.17.177.121/24	ie3070-1	e0b	true
	mgmt1	up/up	172.17.178.120/24	ie3070-1	e1a	true
node1	clus1	up/up	172.17.177.122/24	ie3070-2	e0a	true
	clus2	up/up	172.17.177.123/24	ie3070-2	e0b	true
	mgmt1	up/up	172.17.178.122/24	ie3070-2	e1a	true
node2	clus1	up/up	172.17.177.124/24	ie3070-3	e0a	true
	clus2	up/up	172.17.177.125/24	ie3070-3	e0b	true
	mgmt1	up/up	172.17.178.124/24	ie3070-3	e1a	true
node3	clus1	up/up	172.17.177.126/24	ie3070-4	e0a	true
	clus2	up/up	172.17.177.127/24	ie3070-4	e0b	true
	mgmt1	up/up	172.17.178.126/24	ie3070-4	e1a	true

For each cluster management and node management LIF, verify the following:

- The LIF is up.
- The IP address is configured correctly.

For more information about changing the configuration of a LIF, see the *Clustered Data ONTAP Network Management Guide*.

3. Use the `network port show` command to verify that the cluster, node management, and data ports are assigned correctly.

If the cluster consists of a single node, the node's ports will be assigned to the data and node management roles.

Example

```
cluster1::> network port show
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper
node0	e0a	cluster	up	9000	true/true	full/full	1000/1000
	e0b	cluster	up	9000	true/true	full/full	1000/1000
	e0c	data	up	1500	true/true	full/full	1000/1000
	e0d	data	up	1500	true/true	full/full	1000/1000
	e1a	mgmt	up	1500	true/true	full/full	1000/1000

node1							
e0a	cluster	up	9000	true/true	half/full	10/1000	
e0b	cluster	up	9000	true/true	half/full	10/1000	
e0c	data	up	1500	true/true	half/full	10/1000	
e0d	data	up	1500	true/true	half/full	10/1000	
e1a	mgmt	up	1500	true/true	full/full	1000/1000	
node2							
e0a	cluster	up	9000	true/true	full/full	auto/1000	
e0b	cluster	up	9000	true/true	full/full	auto/1000	
e0c	data	up	1500	true/true	full/full	auto/1000	
e0d	data	up	1500	true/true	full/full	auto/1000	
e1a	mgmt	up	1500	true/true	full/full	auto/1000	
node3							
e0a	cluster	up	9000	true/true	full/full	auto/1000	
e0b	cluster	up	9000	true/true	full/full	auto/1000	
e0c	data	up	1500	true/true	full/full	auto/1000	
e0d	data	up	1500	true/true	full/full	auto/1000	
e1a	mgmt	up	1500	true/true	full/full	auto/1000	

Verify that each port has the correct role assigned for your platform. For more information about default port roles and changing the role assignment for a port, see the *Clustered Data ONTAP Network Management Guide*.

Verifying licensing

You should verify that the correct feature licenses are installed on your system.

About this task

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

Step

1. Use the `system license show` command to verify that the correct feature licenses are installed on your system by verifying license names as listed in the `Description` column of the command output.

Example

```
cluster1::> system license show
Serial Number: 1-80-123456
Owner: cluster1
Package      Type      Description      Expiration
-----
Base         site     Cluster Base License  -
iSCSI       site     iSCSI License      -
CDMI        site     CDMI License        -
```

```

Serial Number: 1-81-000000000000001122334455
Owner: node1
Package           Type      Description           Expiration
-----
CDMI              license  CIFS License         -
SnapRestore      license  SnapRestore License  -
5 entries were displayed.

```

Verifying the high-availability configuration

If the cluster consists of more than one node, you should verify that storage failover is configured for each HA pair. If you have a two-node cluster, then you should also verify that cluster high availability is configured.

About this task

Single node clusters do not use storage failover.

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

Steps

1. Use the `storage failover show` command to verify that storage failover is enabled for each HA pair.

Example

```

cluster1::> storage failover show
                        Takeover
Node      Partner  Possible State
-----
node0    node1    true    Connected to node1
node1    node0    true    Connected to node0
node2    node3    true    Connected to node3
node3    node2    true    Connected to node2
4 entries were displayed.

```

2. If the cluster consists of only two nodes (a single HA pair), then use the `cluster ha show` command to verify that cluster high availability is configured.

Example

```

cluster1::> cluster ha show
High Availability Configured: true

```

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.

Verifying the system time

You should verify that NTP is configured, and that the time is synchronized across the cluster.

About this task

For more information about managing the system time, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

Steps

1. 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
-----
node0
      ntp1.example.com                   max
      ntp2.example.com                   max
node1
      ntp1.example.com                   max
      ntp2.example.com                   max
node2
      ntp1.example.com                   max
      ntp2.example.com                   max
node3
      ntp1.example.com                   max
      ntp2.example.com                   max
```

2. Use the `cluster date show` command to verify that each node has the same date and time.

Example

```
cluster1::> cluster date show
Node   Date                               Timezone
-----
node0   4/6/2013 20:54:38                 GMT
node1   4/6/2013 20:54:38                 GMT
node2   4/6/2013 20:54:38                 GMT
node3   4/6/2013 20:54:38                 GMT
4 entries were displayed.
```

Setting up a Vserver with FlexVol Volumes

You can create and configure Vservers with FlexVol volumes fully to start serving data immediately or with minimal configuration to delegate administration to the Vserver administrator by using the `vserver setup` command.

Before you begin

You must have understood the [requirements and gathered the required information](#) on page 10 before you start the Vserver Setup wizard or any of the subwizards.

About this task

By using the `vserver setup` command, which launches a CLI wizard, you can perform the following tasks:

- Creating and configuring a Vserver fully
- Creating and configuring a Vserver with minimal network configuration
- Configuring existing Vservers
 - Setting up a network interface
 - Provisioning storage by creating volumes
 - Configuring services
 - Configuring protocols

Note: When you select NDMP as one of the protocols for protocol configuration, NDMP is added to the allowed list of protocols of the Vserver. The Vserver setup wizard does not configure the NDMP protocol.

Steps

1. Depending on your requirements, enter the appropriate command:

If you want to...	Enter the following command...
--------------------------	---------------------------------------

Set up a Vserver by using the Vserver Setup wizard

vserver setup

The `vserver setup` command prompts you to create and configure a Vserver in the following sequence:

- a. Create a Vserver
- b. Create data volumes
- c. Create logical interfaces
- d. Configure services
- e. Configure protocols

The following example shows how to set up a Vserver by using the Vserver Setup wizard:

```
cluster1::>vserver setup
Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

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 your answers to previous questions, and
"exit" if you want to quit the Vserver Setup Wizard. Any changes
you made before typing "exit" will be applied.

You can restart the Vserver Setup Wizard by typing "vserver setup". To
accept a default
or omit a question, do not enter a value.

Vserver Setup wizard creates and configures only data Vservers.
If you want to create a Vserver with Infinite Volume use the vserver
create command.

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.
.....
Enter the Vserver name: vs2.example.com
.....
```

If you want to...	Enter the following command...
--------------------------	---------------------------------------

Set up a network interface for an existing Vserver

vserver setup -vserver vserver_name -network true

vserver_name is the name of the Vserver.

The following example shows how to set up a network interface by using the Vserver Setup wizard:

```
cluster1::> vserver setup -vserver vs2.example.com -network true

Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

.....

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi,
ndmp}
[nfs,cifs,fcp,iscsi,ndmp]:
Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcp,iscsi,ndmp

Step 2: Create a logical interface.
You can type "back", "exit", or "help" at any question.

Do you want to create a logical interface? {yes, no} [yes]:
.....
```

If you want to...	Enter the following command...
--------------------------	---------------------------------------

Provision storage by creating volumes on an existing Vserver

```
vserver setup -vserver vserver_name -storage true
```

vserver_name is the name of the Vserver.

The following example shows how to create volumes by using the Vserver Setup wizard:

```
cluster1::> vserver setup -vserver vs2.example.com -storage true
Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.
```

```
.....
```

```
Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.
```

```
Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi,
ndmp}
```

```
[nfs,cifs,fcp,iscsi,ndmp]:
```

```
Choose the Vserver client services to be configured {ldap, nis, dns}:
```

```
Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcp,iscsi,ndmp
```

```
Step 2: Create a data volume
You can type "back", "exit", or "help" at any question.
```

```
Do you want to create a data volume? {yes, no} [yes]:
```

```
.....
```

Note: You cannot set up IPv6 network interface by using the Vserver setup wizard.

If you want to...	Enter the following command...
--------------------------	---------------------------------------

Configure services for an existing Vserver

If you want to configure your Vserver to use `ldap` or `nis` as the name service (`-ns-switch`), you should also include `file` as a name service. Including `file` as a name service enables the Vserver user account authentication through the Vserver's local administrative repository.

```
vserver setup -vserver vserver_name -services ldap,nis,dns
```

`vserver_name` is the name of the Vserver.

The following example shows how to configure services by using the Vserver Setup wizard:

```
cluster1::> vserver setup -vserver vs2.example.com -services ldap

Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

.....

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver data protocols to be configured {nfs, cifs, fcp, iscsi,
ndmp}
[nfs,cifs,fcp,iscsi,ndmp]:
Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fcp,iscsi,ndmp

Step 2: Configure LDAP (Lightweight Directory Access Protocol).
You can type "back", "exit", or "help" at any question.

Do you want to configure LDAP? {yes, no} [yes]:
.....
```

If you want to... **Enter the following command...**

Configure protocols for an existing Vserver

You must have created LIFs for the protocols.

```
vserver setup -vserver vserver_name -protocols
nfs,cifs,iscsi,fc, ndmp
```

vserver_name is the name of the Vserver.

Note: When you specify the protocols value as `fc`, you can configure both FC and FCoE for a Vserver.

The following example shows how to configure protocols by using the Vserver Setup wizard:

```
cluster1::> vserver setup -vserver vs2.example.com -protocols iscsi

Welcome to the Vserver Setup Wizard, which will lead you through
the steps to create a virtual storage server that serves data to clients.

.....

Step 1. Create a Vserver.
You can type "back", "exit", or "help" at any question.

Choose the Vserver client services to be configured {ldap, nis, dns}:

Vserver vs2.example.com's allowed protocol list has been modified to
nfs,cifs,fc,iscsi,ndmp

Step 2: Configure iSCSI.
You can type "back", "exit", or "help" at any question.

Do you want to configure iSCSI? {yes, no} [yes]:
.....
```

2. Follow the prompts to complete the Setup wizard:
 - 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.
3. Use the `vserver show` command to verify the newly created Vserver.

You can view the attributes of the Vserver in detail by using the `vserver show -instance` command.

Example

The following example shows how to display information about all existing Vservers:

```
cluster1::>vserver show
```

Vserver	Type	Admin State	Root Volume	Aggregate	Name Service	Name Mapping
vs1.example.com	data	running	root_voll	aggr1	file	file
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-
cluster1-02	node	-	-	-	-	-

```
vs2.example.com  data  running  root_vol2  aggr2  file  file
5 entries were displayed.
```

Result

When a Vserver is created, its root volume of 1 GB size is created. When you set up a Vserver, it is started automatically and is in running state. By default, the vsadmin user account is created and is in the locked state. The vsadmin role is assigned to the default vsadmin user account.

After you finish

To delegate the administration to a Vserver administrator, you must set up a password, unlock the vsadmin user account, and create a LIF for accessing and enable the firewall policy for managing the Vserver.

For more information about delegating administration to a Vserver administrator, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.

If you want to change the role of the default user vsadmin, you must set the password for vsadmin user before changing the role.

For Vservers with FlexVol volumes, it is best to create a load-sharing mirror copy for the root volume on each node of the cluster so that the namespace directory information remains available in the event of a node outage or failover. For more information about creating load-sharing mirror copy, see the *Clustered Data ONTAP Logical Storage Management Guide*.

Where to go from here

After setting up the software, you can use the NetApp Support Site to find information about how to configure your storage system for third-party storage, provision storage, and manage the cluster.

For information about configuring the software to use third-party storage in a V-Series system, see the *Clustered Data ONTAP Physical Storage Management Guide*.

For information about provisioning your storage, see the *Clustered Data ONTAP Logical Storage Management Guide*.

To find documentation about managing your storage system after the software is set up, see the *Clustered Data ONTAP Documentation Map*.

Related information

The NetApp Support Site: 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

AutoSupport
setup [24](#)

C

cluster
 verifying health after cluster setup [32](#)
 verifying system time is synchronized [39](#)
cluster health
 verifying [32](#)
cluster network
 verifying connectivity after cluster setup [34](#)
cluster replication rings
 verifying cluster is in RDB quorum [33](#)
cluster setup
 creating the cluster [18](#)
 information to gather for [5](#)
 joining a node [19](#)
 preparing for [4](#)
 process for [17](#)
 setting up configuration backup schedules [23](#)
 synchronizing the system time [21](#)
 verifying [32](#)
clusters
 configuring HA in two-node [20](#)
 verifying cluster is in RDB quorum [33](#)
 verifying cluster is in RDB quorum after setup [33](#)
configuration backup schedules
 setting up for single node clusters [23](#)

E

Event Management System
 setting up [26](#)
event messages
 reducing number of [26](#)

F

failover
 testing storage [38](#)
feature licenses
 See licenses

G

giveback
 testing [38](#)

H

HA
 configuring in two-node clusters [20](#)
high availability
 configuring in two-node clusters [20](#)
 verifying the configuration [37](#)

L

licenses
 verifying after cluster setup [36](#)
logical interfaces (LIFs)
 verifying connectivity after cluster setup [34](#)

M

messages
 configuring EMS [26](#)

N

NetApp Support Site
 registering on [5](#)
network
 configuring the SP [29](#)
Network Time Protocol (NTP)
 associating nodes with [21](#)
 enabling for the cluster [21](#)
node
 joining to a cluster [19](#)
nodes
 renaming [30](#)
 verifying participation in RDB quorums [33](#)
NTP
 See Network Time Protocol

P

ports

verifying role assignments after cluster setup [34](#)

Q

quorum

verifying the cluster is in [33](#)

R

replication rings

verifying cluster is in RDB quorum [33](#)

requirements

for software setup [4](#)

S

setup

AutoSupport [24](#)

single node clusters

setting up configuration backup schedules for [23](#)

software configuration

resources for [47](#)

software setup

preparing for [4](#)

requirements for [4](#)

SPs

configuring the network [29](#)

storage failover

verifying the configuration [37](#)

synchronizing

system time across the cluster [21](#)

system time

synchronizing [21](#)

verifying [39](#)

T

takeover

testing [38](#)

testing

storage failover [38](#)

third-party storage

resources for configuring [47](#)

V

verifying

cluster health [32](#)

cluster is in quorum [33](#)

cluster setup [32](#)

high-availability configuration [37](#)

licenses [36](#)

network connectivity after cluster setup [34](#)

storage failover configuration [37](#)

Vserver setup

using the Vserver Setup wizard [40](#)

Vservers

information to gather for [10](#)

requirements [10](#)

Vserver setup worksheet [10](#)