IBM

IBM System Storage N series

Clustered Data ONTAP 8.2 Software Setup Guide

Release Candidate Documentation - Contents Subject To Change

Contents

| Preface | 5 |
|--|----|
| About this guide | 5 |
| Supported features | 5 |
| Websites | 5 |
| Getting information, help, and service | 6 |
| Before you call | 6 |
| Using the documentation | 6 |
| Hardware service and support | 7 |
| Firmware updates | 7 |
| How to send your comments | 7 |
| Preparing for the software setup process | 8 |
| Requirements for software setup | 8 |
| Completing the cluster setup worksheet | 9 |
| Completing the SVM setup worksheet | 14 |
| Setting up the cluster | 21 |
| Creating the cluster on the first node | 22 |
| Joining a node to the cluster | 23 |
| Rebooting each node's Service Processor | |
| Synchronizing the system time across the cluster | 25 |
| Setting up AutoSupport | |
| Setting up the Event Management System | |
| Setting up the Service Processor | 31 |
| Renaming a node | 32 |
| Verifying cluster setup | 34 |
| Verifying cluster health | 34 |
| Verifying that the cluster is in quorum | 35 |
| Verifying network connectivity | 36 |
| Verifying licensing | 39 |
| Verifying the configuration backup schedule for single node clusters | |
| Verifying the high-availability configuration | 41 |
| Testing storage failover | 42 |
| Setting up Storage Encryption | 43 |

| What Storage Encryption is | 43 |
|--|----|
| Limitations of Storage Encryption | 43 |
| Information to collect before configuring Storage Encryption | 44 |
| Using SSL for secure key management communication | 45 |
| Requirements for SSL certificates | 45 |
| Installing SSL certificates on the storage system | 46 |
| Running the Storage Encryption setup wizard | 47 |
| Setting up SVMs with FlexVol Volumes | 49 |
| Where to go from here | 51 |
| Copyright information | 52 |
| Trademark information | 53 |
| Index | 56 |
| | |

Preface

About this guide

This document applies to IBM N series systems running Data ONTAP, including systems with gateway functionality. If the terms *Cluster-Mode* or *clustered Data ONTAP* are used in this document, they refer to the Data ONTAP features and functionality designed for clusters, which are different from 7-Mode and prior Data ONTAP 7.1, 7.2, and 7.3 release families.

In this document, the term *gateway* describes IBM N series storage systems that have been ordered with gateway functionality. Gateways support various types of storage, and they are used with thirdparty disk storage systems—for example, disk storage systems from IBM, HP[®], Hitachi Data Systems[®], and EMC[®]. In this case, disk storage for customer data and the RAID controller functionality is provided by the back-end disk storage system. A gateway might also be used with disk storage expansion units specifically designed for the IBM N series models.

The term *filer* describes IBM N series storage systems that either contain internal disk storage or attach to disk storage expansion units specifically designed for the IBM N series storage systems. Filer storage systems do not support using third-party disk storage systems.

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in *Websites* on page 5).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

• A listing of currently available N series products and features can be found at the following web page:

www.ibm.com/storage/nas/

• The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web

6 | Software Setup Guide

content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:

www.ibm.com/storage/support/nseries/

This web page also provides links to AutoSupport information as well as other important N series product resources.

• IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:

www.ibm.com/systems/storage/network/interophome.html

• For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in *Websites* on page 5) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in *Websites* on page 5).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in *Websites* on page 5).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to *starpubs@us.ibm.com*.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- · A detailed description of any information that should be changed

Preparing for the software setup process

Before setting up the software, you must complete the software setup prerequisites, register on the N series support website (accessed and navigated as described in *Websites* on page 5), and gather cluster and Storage Virtual Machine (SVM) configuration information.

Requirements for software setup

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

Site requirements

Your site must meet the physical, connectivity, power, and model-specific requirements for your cluster. For more information, see the *IBM System Storage N series Introduction and Planning 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 *Gateway Installation Requirements and Reference Guide*, and the *Gateway Implementation Guide for Third-Party Storage*.
- The serial console should be connected to the cluster.

Related information

IBM N series support web site: www.ibm.com/storage/support/nseries/

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 configured to use network switches)

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 that are connected using network switches. Single-node clusters and two-node switchless clusters do not use a cluster network.

| Types of information | Default | Your values |
|--|--|-------------|
| Private cluster network ports | See the <i>Clustered</i> <i>Data ONTAP</i> <i>Network</i> <i>Management</i> <i>Guide</i> . | |
| MTU size for cluster ports 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 | |

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 " ". | |

| Types of information | Your values |
|--|-------------|
| Cluster base license key | |
| To get this license key, contact technical | |
| support. | |

Feature license keys

For information about finding license keys for your initial or add-on software orders, contact technical support.

| Types of information | Your values |
|----------------------|-------------|
| Feature license keys | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |

Types of information Your values Cluster administrator password The password for the admin account that the cluster requires before granting cluster administrator access to the console or through a secure protocol. The default rules for passwords are as follows: ٠ A password must be at least eight characters long. • A password must contain at least one letter and one number. Cluster management interface port The physical port that is connected to the data network and enables the cluster administrator to manage the cluster. 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. Cluster management interface IP address A unique IP address for the cluster management interface. The cluster administrator uses this address to access the admin SVM and manage the cluster. Typically, this address should be on the data network. You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization. Example: 192.0.2.66 Cluster management interface netmask The subnet mask that defines the range of valid IP addresses on the cluster management network. Example: 255.255.255.0 Cluster management interface default gateway The IP address for the router on the cluster management network.

Admin Storage Virtual Machine (SVM)

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

Node information (for each node in the cluster)

| Types of information | Your values |
|---|-------------|
| Physical location of the controller | |
| 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"). | |
| Node management interface port | |
| The physical port that is connected to the node management network and enables the cluster administrator to manage the node. | |
| 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. | |
| Node management interface IP address | |
| 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. | |
| You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization. | |
| Example: 192.0.2.66 | |

| Types of information | Your values |
|--|-------------|
| Node management interface netmask 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. Example: 255.255.255.0 | |
| Node management interface default gateway The IP address for the router on the node management network. | |
| System configuration backup destination address (single-node clusters only) The remote URL where the cluster configuration backups will be uploaded. You can specify either an HTTP or FTP address. Note: The web server that serves the remote URL must have PUT operations enabled. | |
| User name for the configuration backup destination address (single-node clusters only) The user name required to log in to the remote URL and upload the configuration backup file. | |
| Password for the configuration backup destination address (single-node clusters only) The password for the remote URL, if the user name requires a password. | |
| Service Processor (SP) interface IP address* The IP address for the node's SP. | |
| SP interface netmask* The subnet mask that defines the range of valid IP addresses on the SP network. | |
| SP interface default gateway* The IP address of the gateway for the SP network. | |

14 | Software Setup Guide

* The Cluster Setup wizard provides you with the option to enable the SP network to use DHCP. If you intend to use DHCP, you do not need to gather the SP network values; they will be configured automatically using DHCP.

NTP server information

| Types of information | Your values |
|--|-------------|
| NTP server address | |
| 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. | |

Completing the SVM setup worksheet

Before you start the Vserver Setup wizard to create and configure a Storage Virtual Machine (SVM), you must gather the required information to complete the wizard successfully.

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

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

- 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:

SVM information

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

16 | Software Setup Guide

| Types of information | Your values |
|--|-------------|
| SVM root volume's security style | |
| Determines the type of permissions that can be used to control data access to a volume | |
| For more information about the security styles, see the <i>Clustered Data ONTAP File Access</i> <i>Management Guide for CIFS</i> or <i>Clustered Data</i> <i>ONTAP File Access Management Guide for</i> <i>NFS</i> . | |

Information for creating volumes on the SVM

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

Information for creating an IP network interface on the SVM

| Types of information | Values |
|--|--------|
| <i>LIF name</i> The default LIF name is used if you do not specify one. | |
| Protocols Protocols that can use the LIF Note: Protocols that can use the LIF cannot be modified after the LIF is created. | |

| Types of information | Values |
|---|--------|
| Home node | |
| 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. | |
| Home port | |
| 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. | |
| IP address | |
| Network mask | |
| Default gateway IP address | |

Information for creating an FC network interface on the SVM

| Types of information | Values |
|---|--------|
| LIF name | |
| The default LIF name is used if you do not specify one. | |
| Protocols | |
| Protocols that can use the LIF | |
| Note: Protocols that can use the LIF cannot be modified after the LIF is created. | |
| Home node | |
| 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. | |
| Home port | |
| 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. | |

Information for configuring LDAP

| Types of information | Values |
|------------------------|--------|
| LDAP server IP address | |

18 | Software Setup Guide

| Types of information | Values |
|---|--------|
| <i>LDAP server port number</i> The default LDAP server port number is used if you do not specify one. | |
| LDAP server minimum bind authentication level | |
| Bind domain name and password | |
| Base domain name | |

Information for configuring NIS

| Types of information | Values |
|---------------------------------|--------|
| NIS domain name | |
| IP addresses of the NIS servers | |

Information for configuring DNS

| Types of information | Values |
|---------------------------------|--------|
| DNS domain name | |
| IP addresses of the DNS servers | |

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

Information for configuring CIFS protocol

| Types of information | Values |
|---|--------|
| Domain name | |
| <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: ! @ # \$ % & ()_' { } .~ | |

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

Information for configuring iSCSI protocol

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

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

| Types of information | Values |
|---|--------|
| <i>igroup name</i> The default igroup name is used if you do not specify one. | |
| <i>Worldwide port number (WWPN) of the initiators</i> | |
| Operating system type of the initiator | |
| <i>LUN name</i> The default LUN name is used if you do not specify one. | |

20 | Software Setup Guide

| Types of information | Values |
|---------------------------------------|--------|
| Volume for LUN | |
| Volume that is to be used for the LUN | |
| LUN size | |

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 22

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 Storage Virtual Machine (SVM), add feature license keys, and create the node management interface for the first node.

2. Joining a node to the cluster on page 23

After creating a new cluster, you use the Cluster Setup wizard to join each remaining node to the cluster one at a time. The wizard helps you to configure each node's node management interface.

3. Rebooting each node's Service Processor on page 24

After joining each node to the cluster, you should reboot the Service Processor (SP) on each node to ensure that the hardware-assisted takeover feature is configured with the correct IP address for each node's SP.

4. Synchronizing the system time across the cluster on page 25

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

5. Setting up AutoSupport on page 26

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

6. Setting up the Event Management System on page 28

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.

7. Setting up the Service Processor on page 31

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.

8. *Renaming a node* on page 32

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 Storage Virtual Machine (SVM), 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?

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 setting, 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:

You can access the Cluster Setup wizard to change any of the values you entered for the admin SVM or node SVM 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, you use the Cluster Setup wizard to join each remaining node to the cluster one at a time. The wizard helps you to configure each node's node management interface.

Before you begin

The cluster must be created on the first node.

About this task

You can only join one node to the cluster at a time. When you start to join a node to the cluster, you must complete the join, and the node must be part of the cluster before you can start to join the next node.

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:

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

5. Repeat this task for each remaining node.

Rebooting each node's Service Processor

After joining each node to the cluster, you should reboot the Service Processor (SP) on each node to ensure that the hardware-assisted takeover feature is configured with the correct IP address for each node's SP.

Step

1. Reboot the SP for each node:

system node service-processor reboot-sp -node node_name

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.

26 | Software Setup Guide

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 ntpl.example.com
                                    max
cluster1-02 ntp1.example.com
                                   max
cluster1-03 ntpl.example.com
                                   max
cluster1-04 ntp1.example.com
                                    max
cluster1-05 ntpl.example.com
cluster1-06 ntpl.example.com
                                    max
                                    max
6 entries were displayed.
```

Setting up AutoSupport

You can control whether and how AutoSupport information is sent to 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 that 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 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 |
|--------------------|---|
| -to | Up to five comma-separated individual email addresses or distribution lists in your internal support organization that will receive key AutoSupport messages |
| -noteto | 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 |
| -partner-address | 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:

28 | Software Setup Guide

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

- 7. Check the overall configuration using the system node autosupport show commandwith 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 IBM is receiving your AutoSupport messages by checking the email addresses that were specified in the autosupport.to option.

An automated response from the IBM mail handler should have been sent to these email addresses.

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. Display the mail server settings:

```
event config show
```

Example

```
cluster1::> event config show
Mail From: admin@localhost
Mail Server: localhost
```

2. Optional: If necessary, change the mail server settings to meet your requirements:

event config modify -mailserver name -mailfrom email address

Example

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

cluster1::> event config modify -mailserver mailhost.example.com -mailfrom admin@nodel-example.com cluster1::> event config show Mail From: admin@nodel-example.com Mail Server: mailhost.example.com

3. Create the destination for events by using the event destination create command.

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

Example

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

```
cluster1::> event destination create -name test dest -mail
me@example.com
cluster1::> event destination show
                                      Hide
Name Mail Dest. SNMP Dest. Syslog Dest. Params
_____ ____
allevents -
                                   false
false
                  -
asup -
                            -
            ____
                           -
                                      false
criticals -
pager
      _
                                      false
test_dest me@example.com -
                           -
                                      false
traphost -
                                       false
```

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

You should 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:

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

5. 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:

| <pre>cluster1::> event route</pre> | show -sever: | ity <=CRITICAL | Freq | Timo |
|--|----------------------------|-------------------------------------|-------------|----------------|
| Message Threshd | Severity | Destinations | Threshd | TIME |
| | | | | |
| adminapi.time.zoneDiff api.engine.killed app.log.alert | ALERT CRITICAL ALERT | test_dest test_dest test_dest | 0 0 0 | 3600 0 0 |
| app.log.crit app.log.emerg | CRITICAL EMERGENCY | test_dest test_dest | 0 0 | 0 0 |

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

Example

For example, the following event is sent to the destinations no more than once per hour:

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

Result

When you have completed these steps, all critical and above 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 Node | show Health | Eligibility |
|----------------------------------|------------------------------|------------------------------|
| node0 node1 node2 node3 | true true true true | true true true true |
| 4 entries were dispi | ayeu. | |

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 | cluster ring show -unitname mgmt | | | |
| Volume location database | cluster ring show -unitname vldb | | | |
| Virtual-Interface manager | cluster ring show -unitname vifmgr | | | |
| SAN management daemon | cluster ring show -unitname bcomd | | | |

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.1.x:

Example

This example shows the volume location database process for a cluster running Data ONTAP 8.2.x:

Release Candidate Documentation - Contents Subject To Change

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 each node is in a SAN quorum:

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

The most recent scsiblade event message for each node should indicate that the scsiblade is in quorum.

If a node is out of SAN quorum, you can use the storage failover takeover and storage failover giveback commands to perform a planned takeover and giveback with the node's high-availability partner and bring the node back into SAN quorum.

Example

In this example, both nodes in the cluster are in SAN quorum.

```
cluster1::> event log show -messagename scsiblade.*

Time Node Severity Event

2/13/2013 14:03:51 node0 INFORMATIONAL scsiblade.in.quorum: The scsi-blade ...

8/13/2013 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 pingcluster 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 10.254.42.25
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 | | | | | | | |
|------------------------------------|-------------|------------|-------------------|----------|---------|------|--|
| | Logical | Status | Network | Current | Current | Is | |
| Vserver | Interface | Admin/Oper | Address/Mask | Node | Port | Home | |
| | | | | | | | |
| cluster1 | | | | | | | |
| | cluster_mgm | nt up/up | 172.17.178.119/24 | ie3070-1 | ela | 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 | ela | true | |
| nodel | | | | | | | |
| | 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 | ela | 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 | ela | 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 | ela | true | |
| | | | | | | | |

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

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

38 | Software Setup Guide

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 | | | | | | | |
|-------------------------------|------|------------|------|------|------------|------------------|--------------|
| | | | | | Auto-Negot | Duplex | Speed (Mbps) |
| Node | Port | Role | Link | MTU | Admin/Oper | Admin/Oper | Admin/Oper |
| | | | | | | | |
| nodeu | | - . | | | | <pre>c / c</pre> | 1000/1000 |
| | e0a | cluster | up | 9000 | true/true | full/full | 1000/1000 |
| | e0b | cluster | up | 9000 | true/true | tull/tull | 1000/1000 |
| | e0c | data | up | 1500 | true/true | full/full | 1000/1000 |
| | e0d | data | up | 1500 | true/true | full/full | 1000/1000 |
| | ela | mgmt | up | 1500 | true/true | full/full | 1000/1000 |
| nodel | | | | | | | |
| | 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 |
| | ela | 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 |
| | ela | 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 |
| | ela | 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 cluster.

About this task

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

Step

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

Example

Verifying the configuration backup schedule for single node clusters

If the cluster consists of a single node, you should verify that the configuration backup schedule is configured 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.

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. 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.

3. 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
```

4. Return to the admin privilege level:

```
set -privilege admin
```

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 cluster 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.
- 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.

Setting up Storage Encryption

During initial setup, your storage system checks whether it is properly configured with selfencrypting disks and is running a version of Data ONTAP that supports Storage Encryption. If the check is successful, you can then launch the Storage Encryption setup wizard after completion of the storage system setup wizard.

What Storage Encryption is

Storage Encryption is an optional feature that you can enable for additional data protection. It is available on certain supported storage controllers and disk shelves that contain disks with built-in encryption functionality.

In a standard storage environment, data is written to disk in cleartext format. This makes the data vulnerable to potential exposure to unauthorized users when disks removed from a storage system are lost or stolen.

When you enable Storage Encryption, the storage system protects your data at rest by storing it on self-encrypting disks.

The authentication keys used by the self-encrypting disks are stored securely on external key management servers.

Limitations of Storage Encryption

You must keep certain limitations in mind when using Storage Encryption.

- For the latest information about which storage systems, disk shelves, and key management servers are supported with Storage Encryption, see the N series Interoperability Matrices website (accessed and navigated as described in *Websites* on page 5).
- All disks in the storage system and optional attached disk shelves must have encryption functionality to be able to use Storage Encryption. You cannot mix regular non-encrypting disks with self-encrypting disks.
- Storage Encryption is not supported with Flash Pool aggregates.
- Storage Encryption key_manager commands are only available for local nodes. They are not available in takeover mode for partner nodes.
- Do not configure Storage Encryption to use 10 Gigabit network interfaces for communication with key management servers. This limitation does not apply to serving data.
- Storage Encryption supports a maximum of 128 authentication keys per key management server. You receive a warning when the number of stored authentication keys reaches 100. You cannot create new authentication keys when the number of stored authentication keys reaches the limit of 128. You must then delete unused authentication keys before you can create new ones.

Related information

IBM N series Interoperability Matrix: www.ibm.com/systems/storage/network/interophome.html

Information to collect before configuring Storage Encryption

You must gather certain information to successfully set up Storage Encryption on your storage system.

| Information to collect | Details | Required | Optional |
|---|--|----------|----------|
| Network interface name | You must provide the name of the network interface the storage system should use to communicate with external key management servers. Note: Do not configure 10 Gigabit network interfaces for communication with key management servers. | X | |
| Network interface IP address | You must provide the IP address of the network interface. | Х | |
| Network interface subnet mask | You must provide the subnet mask of the network interface. | Х | |
| Network interface gateway IP address | You must provide the IP address for the network interface gateway. | Х | |
| IP address for external key management server | You must link the storage system to at least one external key management server during setup. | Х | |
| IP address for additional external key management servers | You can link the storage system to multiple additional external key management servers during setup for redundancy. | | Х |
| Port number for each external key management server | You must provide the port number that each key management server listens on. The port number must be the same for all key management servers. | Х | |
| Public SSL certificate for storage system | You must provide a public SSL certificate for the storage system to link it to the external key management server. | Х | |
| Private SSL certificate for storage system | You must provide a private SSL certificate for the storage system. | Х | |

| Information to collect | Details | Required | Optional |
|--|---|----------|----------|
| Public SSL certificate for external key management servers | You must provide a public SSL certificate for each external key management server to link it to the storage controller. | Х | |
| Key tag name | You can provide a name that is used to identify all keys belonging to a particular storage system. The default key tag name is the system's host name. | | Х |

Using SSL for secure key management communication

The storage system and key management servers use SSL connections to keep the communication between them secure. This requires you to obtain and install various SSL certificates for the storage system and each key management server before you can set up and configure Storage Encryption.

To avoid issues when installing SSL certificates, you should first synchronize the time between the following systems:

- the server creating the certificates
- the key management servers
- the storage system

Requirements for SSL certificates

Before obtaining and installing SSL certificates, you must understand what certificates are required and their requirements.

SSL certificates for Storage Encryption must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format and follow a strict naming convention. The following table describes the required certificate types and naming conventions:

| Certificate for | Certificate type | Certificate file name |
|--------------------------|---------------------|---|
| Storage system | Public | client.pem |
| Storage system | Private | client_private.pem |
| Key management server | Public | key_management_server_ipaddress_CA.pem key_management_server_ipaddress must be identical to the IP address of the key management server that you use to identify it when running the Storage Encryption setup program. |

These public and private certificates are required for the storage system and key management servers to establish secure SSL connections with each other and verify each other's identities.

46 | Software Setup Guide

The certificates for the storage system are only used by the storage system's KMIP client.

The private certificate can be passphrase protected during creation. In this case, the Storage Encryption setup program prompts you to enter the passphrase.

If your key management server does not accept self-signed certificates, you also need to include the necessary certificate authority (CA) public certificate.

In an HA pair, both nodes must use the same public and private certificates.

If you want multiple HA pairs that are connected to the same key management server to have access to each other's keys, all nodes in all HA pairs must use the same public and private certificates.

Installing SSL certificates on the storage system

You install the necessary SSL certificates on the storage system by using the keymgr install cert command. The SSL certificates are required for securing the communication between the storage system and key management servers.

Before you begin

You must have obtained the public and private certificates for the storage system and the public certificate for the key management server and named them as required.

Steps

1. Access the nodeshell by entering the following command:

```
system node run -node node_name
```

- 2. Copy the certificate files to a temporary location on the storage system.
- **3.** Install the public certificate of the storage system by entering the following command at the storage system prompt:

keymgr install cert /path/client.pem

4. Install the private certificate of the storage system by entering the following command at the storage system prompt:

keymgr install cert /path/client_private.pem

5. Install the public certificate of the key management server by entering the following command at the storage system prompt:

keymgr install cert /path/key_management_server_ipaddress_CA.pem

- **6.** If you are linking multiple key management servers to the storage system, repeat Step 4 for each public certificate of each key management server.
- 7. Exit the nodeshell and return to the clustershell by entering the following command:

exit

Running the Storage Encryption setup wizard

You launch the Storage Encryption setup wizard by using the key_manager setup command. You should run the Storage Encryption setup wizard after you complete setup of the storage system and the storage volumes or when you need to change Storage Encryption settings after initial setup.

Steps

1. Access the nodeshell by entering the following command:

system node run -node *node_nam*e

2. Enter the following command at the storage system prompt:

key_manager setup

- 3. Complete the steps in the wizard to configure Storage Encryption.
- Exit the nodeshell and return to the clustershell by entering the following command: exit

Example

The following command launches the Storage Encryption setup wizard and shows an example of how to configure Storage Encryption:

```
storage-system*> key_manager setup
Found client certificate file client.pem.
Registration successful for client.pem.
Found client private key file client_private.pem.
Is this file protected by a passphrase? [no]:
Registration successful for client_private.pem.
Enter the IP address for a key server, 'q' to quit: 172.22.192.192
Enter the IP address for a key server, 'q' to quit: q
Enter the TCP port number for kmip server [6001] :
```

You will now be prompted to enter a key tag name. The key tag name is used to identify all keys belonging to this Data ONTAP system. The default key tag name is based on the system's hostname.

Would you like to use <storage-system> as the default key tag name? [yes]:

Registering 1 key servers... Found client CA certificate file 172.22.192.192_CA.pem. Registration successful for 172.22.192.192_CA.pem. Registration complete.

You will now be prompted for a subset of your network configuration setup. These parameters will define a pre-boot network environment allowing secure connections to the registered key server(s).

Enter network interface: e0a Enter IP address: 172.16.132.165 Enter netmask: 255.255.252.0 Enter gateway: 172.16.132.1 Do you wish to enter or generate a passphrase for the system's encrypting drives at this time? [yes]: yes Would you like the system to autogenerate a passphrase? [yes]: yes Key ID: 080CDCB2000000000100000000000003FE505B0C5E3E76061EE48E02A29822C Make sure that you keep a copy of your passphrase, key ID, and key taq name in a secure location in case it is ever needed for recovery purposes. Should the system lock all encrypting drives at this time? yes Completed rekey on 4 disks: 4 successes, 0 failures, including 0 unknown key and 0 authentication failures. Completed lock on 4 disks: 4 successes, 0 failures, including 0 unknown key and 0 authentication failures.

Setting up SVMs with FlexVol Volumes

You can create and configure Storage Virtual Machines (SVMs) with FlexVol volumes fully to start serving data immediately or with minimal configuration to delegate administration to the SVM administrator by using the vserver setup command.

Before you begin

You must have understood the *requirements and gathered the required information* on page 14 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 SVMs fully
- Creating and configuring SVMs with minimal network configuration
- Configuring existing SVMs
 - 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 SVM. The Vserver Setup wizard does not configure the NDMP protocol.

When you start the Vserver Setup wizard, the following is displayed:

```
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.
```

50 | Software Setup Guide

Steps

- 1. Use the vserver setup command to launch the wizard and create a fully configured SVM.
- 2. Use the vserver show command to verify the newly created SVM.

You can view the attributes of the SVM in detail by using the vserver show $\mbox{-instance}$ command.

Example

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

| cluster1::>vserver show | | | | | | |
|---------------------------|-------|---------|-----------|-----------|---------|---------|
| Vserver | Туре | State | Volume | Aggregate | Service | Mapping |
| vsl.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 an SVM is created, its root volume of 1 GB size is created. When you set up an SVM, 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 an SVM administrator, you must set up a password, unlock the vsadmin user account, create a LIF for accessing the SVM, and enable the firewall policy for managing the SVM.

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

For SVMs 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 N series support site to find information about how to configure your cluster for third-party storage, how to provision storage, and how to manage the cluster.

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

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

For information about provisioning your storage by using Infinite Volumes, see the *Clustered Data ONTAP Infinite Volumes Management Guide*.

Related information

IBM N series support website: www.ibm.com/storage/support/nseries

Copyright and trademark information

Copyright ©1994 - 2014 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2014 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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.A. 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

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 http://www.ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, 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.

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.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, N.Y. 10504-1785 U.S.A.

For additional information, visit the web at: http://www.ibm.com/ibm/licensing/contact/

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

A

AutoSupport setup 26

С

certificates installing SSL, on storage systems 46 SSL requirements 45 cluster adding nodes to 23 verifying health after cluster setup 34 cluster health verifying 34 cluster network verifying connectivity after cluster setup 36 cluster replication rings verifying cluster is in RDB quorum 35 cluster setup creating the cluster 22 information to gather for 9joining a node 23preparing for 8 process for 21 synchronizing the system time 25 verifying 34 verifying configuration backup schedules 40 clusters verifying cluster is in RDB quorum 35 verifying cluster is in RDB quorum after setup 35 configuration backup schedules verifying for single node clusters 40

E

Event Management System setting up 28 event messages reducing number of 28

F

failover testing storage 42 feature licenses See licenses

G

giveback testing 42

H

high availability verifying the configuration 41

I

installing SSL certificates on storage systems 46

K

key management servers secure communication 45

L

licenses verifying after cluster setup *39* limitations Storage Encryption *43* logical interfaces (LIFs) verifying connectivity after cluster setup *36*

Μ

messages configuring EMS 28

N

network configuring the SP *31* Network Time Protocol (NTP) associating nodes with *25* enabling for the cluster *25*

57 | Software Setup Guide

nodes joining to a cluster 2.3 rebooting the SP 24 renaming 32 verifying participation in RDB quorums 35 NTP See Network Time Protocol

P

verifying role assignments after cluster setup 36

Q

quorum verifying the cluster is in 35

R

replication rings verifying cluster is in RDB quorum 35 requirements for software setup 8

S

secure communication between storage systems and key management servers 45 setting up Storage Encryption 43 setup AutoSupport 26 setup wizards running the Storage Encryption 47 single node clusters verifying configuration backup schedules for 40 software configuration resources for 51 software setup preparing for 8requirements for 8SPs configuring the network 31 rebooting 24 SSL certificates installing on storage systems 46 requirements 45

SSL connections for secure communication 45Storage Encryption explained 43 information to collect before configuring 44 installing SSL certificates for 46 limitations 43 running the setup wizard 47secure communication 45 setting up 43 SSL certificates requirements 45 storage failover verifying the configuration 41 **SVMs** information to gather for 14 requirements 14 Vserver setup worksheet 14 SVMs setup using the Vserver Setup wizard 49 synchronizing system time across the cluster 25 system time synchronizing 25

Т

takeover testing 42 testing storage failover 42 third-party storage resources for configuring 51

V

verifying cluster health 34 cluster is in quorum 35 cluster setup 34 high-availability configuration 41 licenses 39 network connectivity after cluster setup 36 storage failover configuration 41

W

wizards running the Storage Encryption setup 47

IBM.®

NA 210-06388_A0, Printed in USA

SC27-6282-01

