

NetApp Storage Deployment Guide



● ● ● SMART BUSINESS ARCHITECTURE

February 2013 Series

Preface

Who Should Read This Guide

This Cisco® Smart Business Architecture (SBA) guide is for people who fill a variety of roles:

- Systems engineers who need standard procedures for implementing solutions
- Project managers who create statements of work for Cisco SBA implementations
- Sales partners who sell new technology or who create implementation documentation
- Trainers who need material for classroom instruction or on-the-job training

In general, you can also use Cisco SBA guides to improve consistency among engineers and deployments, as well as to improve scoping and costing of deployment jobs.

Release Series

Cisco strives to update and enhance SBA guides on a regular basis. As we develop a series of SBA guides, we test them together, as a complete system. To ensure the mutual compatibility of designs in Cisco SBA guides, you should use guides that belong to the same series.

The Release Notes for a series provides a summary of additions and changes made in the series.

All Cisco SBA guides include the series name on the cover and at the bottom left of each page. We name the series for the month and year that we release them, as follows:

month year Series

For example, the series of guides that we released in February 2013 is the “February Series”.

You can find the most recent series of SBA guides at the following sites:

Customer access: <http://www.cisco.com/go/sba>

Partner access: <http://www.cisco.com/go/sbachannel>

How to Read Commands

Many Cisco SBA guides provide specific details about how to configure Cisco network devices that run Cisco IOS, Cisco NX-OS, or other operating systems that you configure at a command-line interface (CLI). This section describes the conventions used to specify commands that you must enter.

Commands to enter at a CLI appear as follows:

```
configure terminal
```

Commands that specify a value for a variable appear as follows:

```
ntp server 10.10.48.17
```

Commands with variables that you must define appear as follows:

```
class-map [highest class name]
```

Commands shown in an interactive example, such as a script or when the command prompt is included, appear as follows:

```
Router# enable
```

Long commands that line wrap are underlined. Enter them as one command:

```
wrr-queue random-detect max-threshold 1 100 100 100 100 100  
100 100 100
```

Noteworthy parts of system output or device configuration files appear highlighted, as follows:

```
interface Vlan64  
ip address 10.5.204.5 255.255.255.0
```

Comments and Questions

If you would like to comment on a guide or ask questions, please use the [SBA feedback form](#).

If you would like to be notified when new comments are posted, an RSS feed is available from the SBA customer and partner pages.

Table of Contents

What's In This SBA Guide.....	1
Cisco SBA Data Center	1
Route to Success	1
About This Guide	1
Introduction.....	2
Related Reading	2
Business Overview.....	2
Technology Overview.....	3
Deploying a NetApp Storage Solution.....	7
Completing Storage Array Setup.....	7
Provisioning Storage.....	12
Deploying FCoE or iSCSI Storage.....	16
Adding a LUN	22

Increasing Efficiency and Flexibility with Advanced Features.....	27
Thin Provisioning	27
NetApp Deduplication and Compression	27
NetApp Snapshot	27
NetApp FlexClone	27
Backup, Disaster Recovery, and High Availability	28
NetApp Advanced Solutions	29
Monitoring and Management	29
Appendix A: Changes.....	31

What's In This SBA Guide

Cisco SBA Data Center

Cisco SBA helps you design and quickly deploy a full-service business network. A Cisco SBA deployment is prescriptive, out-of-the-box, scalable, and flexible.

Cisco SBA incorporates LAN, WAN, wireless, security, data center, application optimization, and unified communication technologies—tested together as a complete system. This component-level approach simplifies system integration of multiple technologies, allowing you to select solutions that solve your organization's problems—without worrying about the technical complexity.

Cisco SBA Data Center is a comprehensive design that scales from a server room to a data center for networks with up to 10,000 connected users. This design incorporates compute resources, security, application resiliency, and virtualization.

Route to Success

To ensure your success when implementing the designs in this guide, you should first read any guides that this guide depends upon—shown to the left of this guide on the route below. As you read this guide, specific prerequisites are cited where they are applicable.

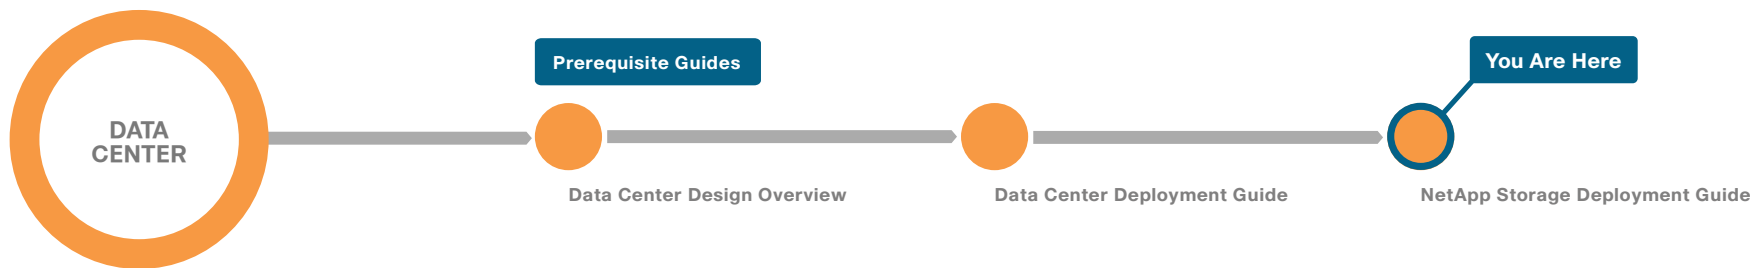
About This Guide

This *ecosystem partner guide* presents solutions, products, or services—provided by a Cisco SBA ecosystem partner—that are compatible with and complementary to SBA.

You can find the most recent series of Cisco SBA guides at the following sites:

Customer access: <http://www.cisco.com/go/sba>

Partner access: <http://www.cisco.com/go/sbachannel>



Introduction

The requirements of smaller organizations increasingly reflect those of larger enterprises, although generally on a smaller scale. However, with limited IT skill sets, many organizations must rely on partners and vendors that can deliver effective solutions in a simplified manner. Cisco SBA solutions address these needs.

NetApp storage solutions offer the performance and functionality to unlock the value of your business, regardless of size. The NetApp family of storage solutions offers class-leading storage efficiency and performance that scale with your organization. Whether you deploy a NetApp fabric attached storage (FAS) system or choose to extend the life of your existing Fibre Channel (FC) storage with a NetApp V-Series system, NetApp storage solutions offer a common storage platform that is designed to maximize the efficiency of your data storage. Capable of data-in-place upgrades to more powerful FAS or V-Series systems, running the same Data ONTAP operating system and using the same management tools and feature sets, NetApp systems grow with your organization.

This guide focuses on deploying a NetApp storage system into networks built upon the Smart Business Architecture. Components of the entire storage solution are present in all three layers of Cisco SBA. The Cisco Ethernet or FC-switching infrastructure forms the network foundation for access to storage, along with providing additional network services, such as security. The NetApp storage systems themselves are end-nodes to the network, but provide additional security services and multiprotocol storage access as a user service, using either block-based access methods, such as Internet Small Computer System Interface (iSCSI), FC, and Fibre Channel over Ethernet (FCoE), or network attached storage (NAS) protocols.

Related Reading

The *Cisco SBA—Data Center Design Overview* provides an overview of the data-center architecture. This guide discusses how Cisco SBA data center architecture is built in layers—the foundation of Ethernet and storage networks and computing resources; the data-center services of security, application resilience, and virtual switching; and the user services layer that contains applications and user services.

The *Cisco SBA—Data Center Deployment Guide* focuses on the processes and procedures necessary to deploy your data-center foundation Ethernet and storage transport. The data-center foundation is designed to support the flexibility and scalability of the Cisco Unified Computing System and provides details for the integration of functionality between the server and the network for Cisco and non-Cisco servers. The foundation design includes data-center services, such as security with firewall and intrusion prevention, and application resiliency with advanced server load-balancing techniques. This guide also discusses the considerations and options for data-center power and cooling. The supplemental *Cisco SBA—Data Center Configuration Files Guide* provides snapshots of the actual platform configurations used in the design.

The *Cisco SBA—Data Center Unified Computing System Deployment Guide* provides the processes and procedures necessary to deploy a Cisco Unified Computing System using both the Cisco B-Series blade server system and Cisco C-Series rack-mount servers to a point where they are ready to deploy an operating system or hypervisor software.

The supplemental *Cisco SBA—Data Center Virtualization with UCS, Nexus 1000V, and VMware Deployment Guide* provides a concise yet detailed process of deploying VMware in your data center and on the Cisco UCS B-Series and C-Series servers. Additionally, the guide details the deployment of the Cisco Nexus 1000V virtual switch to enhance the management and control of the network in a VMware environment.

Business Overview

As the value and amount of electronic data increases, deploying the right storage strategy becomes even more critical. A centralized storage strategy offers the best solution to store, protect, and retain your valuable data. Today's solutions address technology trends around managed scalability, server virtualization, and networking with advanced features and simplified management for organizations with limited IT staff.

Trying to keep up with data growth can be a challenge, especially when using direct attached storage (DAS). Selecting the right amount of storage for a given application is not an exact science. You must estimate how much data you need to store over time, and then choose a server with enough drive bays to accommodate the expected data growth. When the drive bays have all been filled, you can directly attach an external disk shelf or Just a Bunch of Disks (JBOD) via a Small Computer System Interface (SCSI) or Serial-Attached SCSI (SAS), but that is an expensive solution if you expect to use only a portion of the capacity of the shelf. A DAS shelf typically cannot be used by more than one server at a time.

One of the most compelling technology trends today is server virtualization. Server virtualization offers the ability to increase the utilization of excess CPU-processing capacity available with today's high-performing multicore CPUs. Hosting many virtual servers on a physical server CPU allows better utilization of this untapped processing capacity. After a server is virtualized from the physical hardware, it can be moved from one physical server to another to support load balancing, hardware servicing or failover, and site-to-site mobility. These advanced virtualization features are only available in shared networked storage environments.

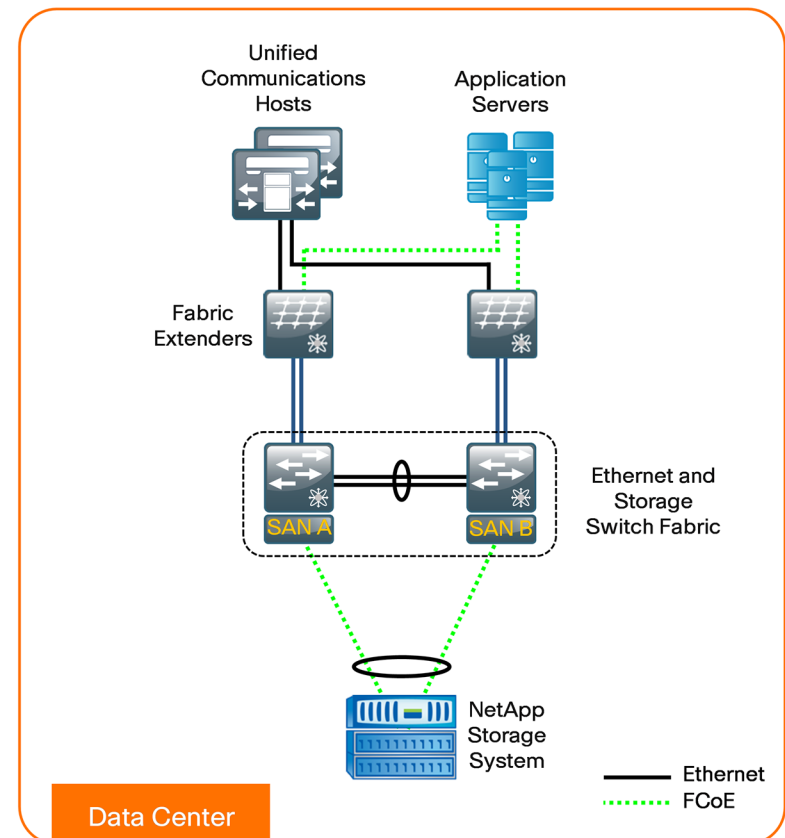
Server virtualization is also driving the need for more robust data networks. The ratio of virtual servers to physical servers is typically 10:1 or greater, resulting in much higher I/O loads per physical server. Ethernet is emerging as an increasingly robust and capable storage network that offers flexibility, simplicity, and performance for organizations of all sizes. The ability to share a common network technology for data traffic, voice, video, and storage offers tremendous value over alternative dedicated solutions. The skill sets required for Ethernet management are common in the industry, and the penetration of this technology into every home and business drives economies-of-scale to keep costs low.

The growth of data is not unique to larger enterprises. Businesses of all sizes face the challenges of acquiring, storing, and retaining large amounts of data. In response to these challenges, solutions that were once considered overly complicated and expensive for smaller businesses have been made simpler to procure, deploy, and manage. Advanced functionality once reserved for costly solutions is now common in lower-priced solutions, making it easier for organizations to address data growth and retention.

Technology Overview

NetApp leads the industry in storage efficiency and innovation. With features such as thin provisioning, deduplication, high-performing and space-efficient NetApp Snapshot copies, and advanced disaster recovery capabilities, NetApp offers solutions to help you achieve your business goals. With a shared vision of a virtual dynamic data center and a partnership that extends back to 2003, NetApp and Cisco are developing technologies that deliver the value and performance to meet your IT requirements. Cisco partners with industry leaders to provide reference architectures and proven configurations to enable organizations of all sizes to meet the IT and business needs of today and tomorrow.

Figure 1 - Storage connectivity over Ethernet



NetApp storage solutions are built on a single platform that scales from small deployments of a few terabytes to large deployments beyond a petabyte, all with a common set of features and management tools. Each NetApp FAS system is capable of running multiple block-based and file-based protocols at the same time, including Network File System (NFS), Common Internet File System (CIFS), FC, FCoE, and iSCSI. NetApp unified storage simplifies data management with the ability to scale your storage environment as your business grows, without the need for staff retraining or forklift equipment upgrades. NetApp V-Series storage systems can extend the life of your existing FC storage investments by extending many of the same advanced features as the FAS storage systems to the management of legacy installed-base storage systems from third-party manufacturers. Figure 1 illustrates IP/Ethernet-based and FCoE-based connections to a NetApp storage system.

Data centers run multiple parallel networks to accommodate both data and storage traffic. To support these different networks in the data center, administrators deploy separate network infrastructures, including different types of host adapters, connectors and cables, and fabric switches. The use of separate infrastructures increases both capital and operational costs for IT executives. The deployment of a parallel storage network, for example, adds to the overall capital expense in the data center, while the incremental hardware components require additional power and cooling, management, and rack space that negatively affect the operational expense.

Consolidating SAN and LAN in the data center into a unified, integrated infrastructure is referred to as network convergence. A converged network reduces both the overall capital expenditure required for network deployment, and the operational expenditure for maintaining the infrastructure.

NetApp Unified Connect delivers data access using NFS, CIFS, iSCSI, and FCoE protocols concurrently over a shared network port using the NetApp unified target adapter. As a leader in Ethernet storage, first as a NAS pioneer and next as an early proponent of iSCSI, NetApp now leads with FCoE and Unified Connect.

The new, industry-leading, NetApp FAS3200 series cost-effectively meets the storage needs of business applications in both virtual and traditional environments. The NetApp FAS3200 family scales to nearly 3-petabytes of versatile storage, which adapts readily to the growing storage demands with even better value and performance than before.

Storage Networking with Cisco SBA

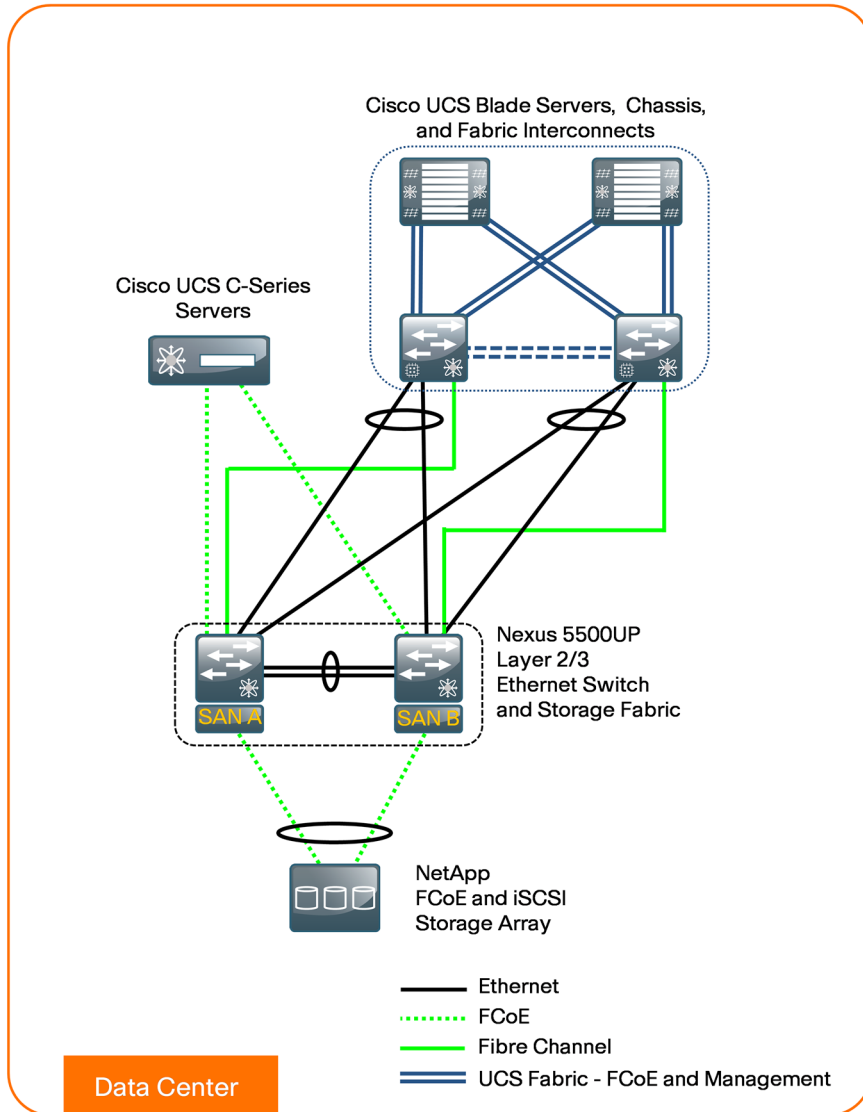
A network infrastructure that uses FC or Ethernet should have no single point of failure. A highly available solution includes:

- Two or more FC or Ethernet network switches.
- Two or more host bus adapters (HBAs) or network interface cards (NICs) per server.
- Two or more target FC ports or Ethernet NICs per storage controller.

When using fibre channel, two fabrics are required to have a truly redundant architecture.

Cisco SBA is designed to address the common requirements of organizations with 250 to 10,000 connected users. Each organization is unique, however, and so are its requirements, so Cisco SBA was built so that additional capabilities could be added without redesigning the network. The Cisco SBA data center foundation provides resilient Ethernet and storage-transport capable of transporting iSCSI, FC, FCoE, and NAS protocols.

Figure 2 - Cisco SBA data center



The Cisco SBA data center foundation is designed to support the flexibility and scalability of the Cisco Unified Computing System and provides details for the integration of functionality between the server, storage, and the network for Cisco and non-Cisco servers.

NetApp has improved upon its leadership position in unified storage as the first storage vendor to support FCoE. FCoE combines two leading technologies—the FC protocol and an enhanced 10-Gigabit Ethernet physical transport—to provide you with more options for SAN connectivity and networking. FCoE allows you to use the same tools and techniques that you use today to manage your FC network and storage. The FCoE network infrastructure offers connectivity to either native NetApp FCoE systems and/or NetApp FC storage systems. This allows you to migrate to a unified Ethernet infrastructure while preserving investments you made in FC storage. FCoE is a logical progression of NetApp's unified storage approach to offering concurrent support for FC, iSCSI, and NAS data access in its enterprise systems. It provides an evolutionary path for FC SAN customers to migrate to Ethernet over time.

These solutions:

- Support SAN (FC, FCoE, and iSCSI) or NAS.
- Scale non-disruptively from a few terabytes to over 3-petabyte.
- Are easily installed, configured, managed, and maintained.
- Dynamically expand and contract storage volumes, as needed.
- Offer features that provide:
 - Rapid backup and recovery with zero-penalty Snapshot copies.
 - Simple, cost-effective replication for disaster recovery.
 - Instant, space-efficient data clones for provisioning and testing.
 - Data deduplication to reduce capacity requirements.

NetApp unified storage solutions enable powerful thin provisioning, simplified data management, and scalable and consistent I/O performance for all protocols across network-attached storage (NAS; NFS, CIFS) and storage networks (SANs; FC, FCoE, and iSCSI) in a single pool of storage.

NetApp storage solutions offer powerful data management and data protection capabilities that enable you to lower costs while meeting capacity, utilization, and performance requirements.

Data Protection

Any consolidation effort increases risk to the organization in the event that the consolidation platform fails. As physical servers are converted to virtual machines (VMs) and multiple VMs are consolidated onto a single physical platform, the effect of a failure to the single platform can be catastrophic. Fortunately, hypervisors provide multiple technologies that enhance the availability of a virtual data center. These technologies include physical server clustering, application load balancing, and the ability to non-disruptively move running VMs and data sets between physical servers.

To ensure storage availability, many levels of redundancy are available for deployments, including purchasing physical servers with multiple storage interconnects or HBAs, deploying redundant storage networking and network paths, and leveraging storage arrays with redundant controllers. A deployed storage design that meets all of these criteria can be considered to have eliminated all single points of failure.

The reality is that data protection requirements in a virtual infrastructure are greater than those in a traditional physical server infrastructure. Data protection is a paramount feature of shared storage devices. NetApp RAID-DP is an advanced Redundant Array of Independent Disks (RAID) technology that is provided as the default RAID level on all NetApp FAS systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy; the overhead with default RAID groups is a mere 12.5 percent. This level of resiliency and storage efficiency makes data residing on RAID-DP safer than data stored on RAID 5 and more cost-effective than RAID 10.

FCOE Target and Unified Connect Requirements

Starting with NetApp Data ONTAP 7.3.2, FCoE is available through unified target adapters (UTAs) in the Data ONTAP 7.3 family. A UTA is more commonly known as a converged network adapter (CNA). In Data ONTAP 8.0.1 and later, FCoE and all other Ethernet protocols normally available from NetApp storage (CIFS, NFS, iSCSI, and so on) are supported concurrently. The FC protocol license is required for FCoE functionality. For other protocols, the relevant license is required. When a UTA is installed in a NetApp storage controller running Data ONTAP 8.0.1 and later releases, a dual-port 10 Gbps NIC and a 10 Gbps FCoE target adapter are presented for each UTA installed in the controller. The UTA is supported on all current platforms that have a PCI-express (PCI-E) expansion slot.

Notes

Deploying a NetApp Storage Solution

By using the NetApp OnCommand System Manager, you can easily configure a NetApp storage system for end-to-end FCoE or iSCSI deployment. System Manager Version 2.0 is supported on all NetApp FAS2000, FAS3000, and FAS6000 systems, and on the corresponding V-Series systems.

The NetApp OnCommand System Manager is a feature-rich, yet easy-to-use, storage management tool for basic configuration and management of NetApp storage systems. System Manager is ideal for initial setup and for configuration of one system at a time. As your environment grows, NetApp offers a suite of storage management tools.

NetApp OnCommand System Manager is the simple yet powerful management tool for NetApp storage. It is easy to use for small to medium-sized businesses, and efficient for large enterprises and service providers.

Process

Completing Storage Array Setup

1. Complete the initial system setup
2. Install OnCommand System Manager
3. Get started with array management

To follow the procedures in this guide, your storage system must be running any of the following versions of Data ONTAP. This guide reflects the use of Data ONTAP 8.0.3 operating in 7-Mode, and NetApp OnCommand System Manager version 2.0.2.

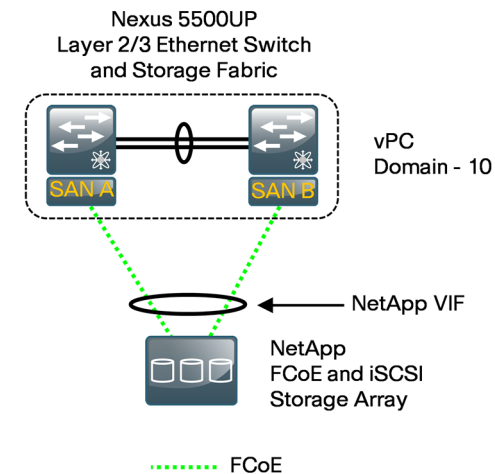
- Data ONTAP 7.x (starting from 7.2.3)
- Data ONTAP 8.x 7-Mode

For system requirements to run NetApp OnCommand System Manager 2.0.2, see the *NetApp System Manager 2.0.2 Quick Start Guide* here: https://library.netapp.com/ecm/ecm_get_file/ECMP1140472

To configure a system that supports Ethernet-based access to storage, it is recommended that you dedicate the e0M interface for system management traffic.

The Cisco SBA data-center design illustrates the storage system dual-attached to the Ethernet switching fabric using a port-channel connection. On the NetApp controller, you will create a virtual interface (VIF), which includes both 10-Gbs Ethernet links from the controller to the data-center core Cisco Nexus 5500UP switches. Configuration for the data-center core Cisco Nexus 5500UP switches is covered in Procedure 1, in the “Deploying FCoE or iSCSI Storage” section of this guide.

Figure 3 - VIF configuration



With a VIF configuration, the dual interfaces provide system resiliency, and when they are configured as a port-channel using link aggregation control protocol (LACP), both links can actively carry storage traffic. In this design, you enable Ethernet storage interfaces for FCoE or iSCSI as a VIF.

Configuration of the FCoE for the unified target adapter in Data ONTAP is the same as it is for a traditional FC target adapter.

Procedure 1 Complete the initial system setup

The setup includes the basic configuration such as host name, IP address, subnet, credentials, management interface, and system default IP gateway.

You can assign a static IP-address to the storage controller manually, or, if you have Dynamic Host Configuration Protocol (DHCP) running on the data-center subnet where your storage controller is installed, you can have DHCP provide the IP address. This procedure uses the static IP-address assignment method using the system setup script.

To configure the NetApp storage system, you need:

- An IP address for system management.
- An IP address for iSCSI traffic, in a different subnet than the system management IP address.
- An IP subnet default gateway address for the iSCSI subnet.

In this procedure, you assign the IP default gateway for the controller to the iSCSI subnet to ensure that iSCSI traffic uses the 10-Gigabit Ethernet interfaces, and not the lower-speed e0M management interface.



Reader Tip

For more information about management network interface configuration, see the NetApp library page here:
<https://library.netapp.com/ecmdocs/ECMM1277792/html/nag/frameset.html>

Step 1: Connect a console to the serial monitor port on the NetApp storage controller and power on your storage system components according to the instructions in the *Installation and Setup Instructions* for your hardware platform.

Step 2: Press **Ctrl+C**. This skips the DHCP IP address search. Next follow the setup script.

Step 3: Assign the storage controller's host name.

A valid hostname consists of alphanumeric characters [a-zA-Z0-9] and dash [-].

Please enter the new hostname []: **NetAppx**

Next, you create the VIF that maps your 10-Gb Ethernet ports for EtherChannel to the data center core for iSCSI and FCoE transport.

Step 4: Enter **y**, enter **1** for a single interface group, and then enter a name for the interface group. This creates an interface group.

Do you want to configure interface groups? [n]: **y**

Number of interface groups to configure? [0] **1**

Name of interface group #1 []: **SBA-VIF-1**

Step 5: Enter **l** for an LACP interface group, enter **i** for load balancing, and then assign two 10-Gb Ethernet links on your controller to the VIF.

Is SBA-VIF-1 a single [s], multi [m] or a lacp [l] interface group? [m] **l**

Is SBA-VIF-1 to use IP-based [i], MAC-based [m], Round-robin based [r] or Port based [p] load balancing? [i] **i**

Number of links for SBA-VIF-1? [0] **2**

Name of link #1 for SBA-VIF-1 []: **e1a**

Name of link #2 for SBA-VIF-1 []: **e1b**

Step 6: Enter the IP address and subnet mask for iSCSI transport on the VIF, and press **Enter** to skip assigning IP addresses to the physical interfaces.

Please enter the IP address for Network Interface SBA-VIF-1 []: **10.4.62.10**

Please enter the netmask for Network Interface SBA-VIF-1 [255.0.0.0]: **255.255.255.0**

Should interface group SBA-VIF-1 take over a partner interface group during failover? (You would answer yes here if you are running redundant controllers and then enter in the second controllers interface name. These names would be identical on each controller) [n] : **n**

Please enter the IP address for Network Interface e0a []:
Should interface e0a take over a partner IP address during failover? [n]: **n**

```
Please enter the IP address for Network Interface e0b []:  
Should interface e0b take over a partner IP address during  
failover? [n]: n
```

Step 7: Enter the IP address and subnet for system management on interface e0M. This must be on a different IP subnet than your iSCSI traffic or you may experience performance issues.

```
e0M is a Data ONTAP dedicated management port.  
NOTE: Dedicated management ports cannot be used for data  
protocols (NFS, CIFS, iSCSI, NDMP or Snap*), and if they are  
configured they should be on an isolated management LAN.  
The default route will use dedicated mgmt ports only as the  
last resort, since data protocol traffic will be blocked by  
default.  
Please enter the IP address for Network Interface e0M  
[]:10.4.63.123  
Please enter the netmask for Network Interface e0M  
[255.0.0.0]: 255.255.255.0  
Should interface e0m take over a partner IP address during  
failover? [n]: n
```

Step 8: Enter **none** for flow control on the e0M management interface, and then enter **n** to continue the setup script.

```
Please enter flow control for e0M {none, receive, send, full}  
[full]: none  
Would you like to continue setup through the web interface?  
[n]: n
```

Step 9: Enter IP address of the default gateway for the controller using the default gateway for the iSCSI IP subnet. By using the iSCSI subnet IP default gateway, you ensure that iSCSI traffic always uses the high-speed 10-Gb Ethernet interfaces configured as a VIF.

```
Please enter the name or IP address of the default gateway:  
10.4.62.1
```

Step 10: Press **Enter** to skip assigning an administration host.

```
The administration host is given root access to the filer's /  
etc files for system administration. To allow /etc root access  
to all NFS clients enter RETURN below.  
Please enter the name or IP address of the administration  
host:
```

Step 11: Enter the time zone and the physical location for this device.

```
Please enter timezone [GMT]: PST8PDT  
Where is the filer located? []: Room 23
```

Step 12: Enter **y** to enable domain name server (DNS) IP-address resolution, enter your DNS domain name, and then enter the IP address of your DNS host. If you have more than one DNS host, enter the IP addresses now.

```
Do you want to run DNS resolver? [n]: y  
Please enter DNS domain name []: cisco.local  
You may enter up to 3 nameservers  
Please enter the IP address for first nameserver []:  
10.4.48.10  
Do you want another nameserver? [n]: n
```

Step 13: Enter **n** to bypass enabling the network information service (NIS), and then press **Enter**.

```
Do you want to run NIS client? [n]: n  
This system will send event messages and weekly reports  
to NetApp Technical Support. To disable this feature,  
enter "options autosupport.support.enable off" within 24  
hours. Enabling AutoSupport can significantly speed problem  
determination and resolution should a problem occur on your  
system. For further information on AutoSupport, please see:  
http://now.netapp.com/autosupport/  
Press the return key to continue.
```

Step 14: Enter **n** to bypass the remote LAN module setup, and then enter **n** to bypass the shelf alternate control path management setup.

```
Would you like to configure the SP LAN interface [y]: n
```


Step 15: Enter a password for the administrative (root) access to your storage controller.

The initial aggregate currently contains 3 disks; you may add more disks to it later using the “aggr add” command.

Setting the administrative (root) password for NetAppx ...

New password: [xxxx]

Retype new password: [xxxx]

Step 16: Log in to the NetApp controller by entering the password that you just configured.

Password: [xxxx]

NetAppx>

Step 17: When the setup is complete, enter **reboot**. This transfers the information that you have entered to the storage system.

NetAppx>**reboot**



Tech Tip

If you do not reboot the system, the information you entered may not be saved.

Step 18: After the system has rebooted, enter the following command to set the storage controller time to the actual time.

NetAppx>**date** [CCyyymmddhhmm]

Where CCyy is the four-digit year, mm is the two-digit month, dd is the two-digit day of the month, hh is the two-digit hour, and the second mm is the two-digit minute.

Step 19: Configure the storage controller to synchronize time with a Network Time Protocol (NTP) server.

NetAppx>**options timed.proto ntp**

Step 20: Add the NTP server to the storage system list.

NetAppx>**options timed.servers 10.4.48.17**

Step 21: Enable NTP synchronization on the storage controller.

NetAppx>**options timed.enable on**

Step 22: Enable Cisco Discovery Protocol (CDP) on the storage controller. This allows the storage controller to identify connectivity over the Ethernet connection to the data center core Nexus 5500 switches.

NetAppx>**options cdpd.enable on**

Step 23: Configure Simple Network Management Protocol (SNMP) communities on the storage controller so that OnCommand System Manager can be queried by your SNMP management system.

NetAppx>**snmp community delete all**

NetAppx>**snmp community add ro cisco**

Step 24: Reinitialize SNMP on the storage controller.

NetAppx>**snmp init 1**

Step 25: Configure the management interface e0M to be used only for management traffic. It is a best practice to set the option interface.blocked.mgmt_data_traffic to on, so that this lower speed interface is not used for storage transport such as iSCSI. This is the default setting on storage systems shipping with Data ONTAP 8.0 and newer.

NetAppx> **options interface.blocked.mgmt_data_traffic on**

Procedure 2

Install OnCommand System Manager

This guide uses the OnCommand System Manager to complete the storage system setup.

Step 1: Install NetApp OnCommand System Manager 2.0 from the software CD, and then launch System Manager.

Step 2: Click **Add**. The Add a System window appears.

Step 3: In the Add a System window, enter the IP Address assigned to the management port e0M in Step 7 of Procedure 1, “Complete the initial system setup”, and then click the arrow next to **More**.

Step 4: In the User Name box, enter **root**.

Step 5: In the **Password** box, enter the password that you configured in Step 15 of Procedure 1, "Complete the initial system setup," and then click **Add**.

The home page is updated with the new controller with the IP address that you statically assigned via CLI.

Step 6: On the Home page, select the new controller IP address, and then on the top bar, click **Login**. The system logs you in using the username and password that you just entered.

Storage system name	Address	Status	Type	Version	Model	System ID
NetApp	10.4.63.123	Up	Storage Controller	8.0.2P4 7-Mode	FAS3240	1579017538

After you log in, the new tab displays the host name and IP address of the controller.

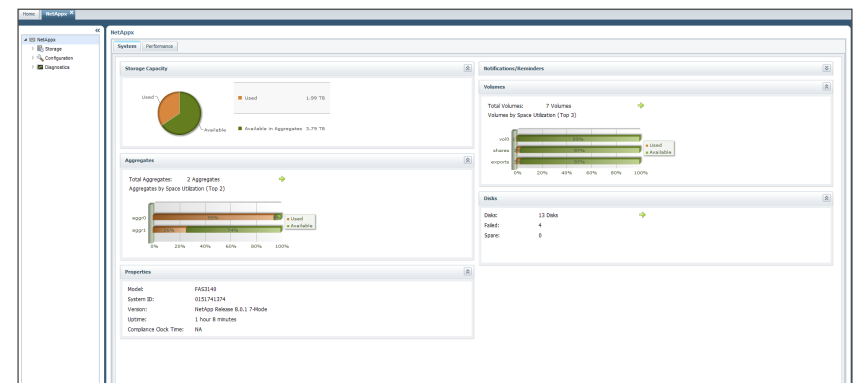
Procedure 3

Get started with array management

Step 1: Click the controller name in the left pane. This allows you to view system details. The dashboard view displays Storage Capacity, Notifications/Reminders, and Properties, including name, IP address, model, system ID, Data ONTAP version, system uptime, and number of volumes, aggregates, and disks.

Step 2: Click the green arrows to the right of the items you want to manage.

System Manager provides basic performance graphs for CPU utilization, total I/O, combined operations of all protocols, and latency for all protocols.



Step 3: System Manager also offers Notifications and Reminders roll-ups. The Notifications area shows the current event list by scanning the syslog. The Reminders area shows reminders or a to-do list. If you want to go directly to the source of the reminder, click the green arrow next to it.

Your NetApp storage system is configured and ready to use. You can now perform storage management (manage disks, aggregates, volumes, qtrees, logical unit numbers (LUNs)), protocol management (CIFS, NFS, iSCSI, FC, FCoE), and system configuration (network, licenses, SNMP, users, groups).

Process

Provisioning Storage

1. Add SAN protocol licenses
2. Create an aggregate
3. Configure flexible volumes

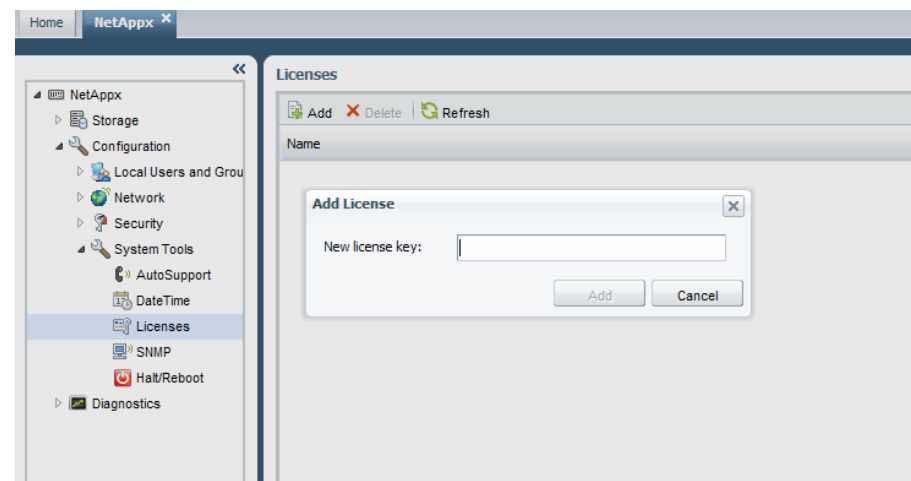
Storage provisioning on NetApp storage is easy, involving only a few steps to provision a LUN or file share. Aggregates form the foundation storage layer from which flexible volumes and then LUNs are stored. The layers of storage virtualization offer a number of advantages to manage and optimize the storage, protection, and retention of your data.

Procedure 1 Add SAN protocol licenses

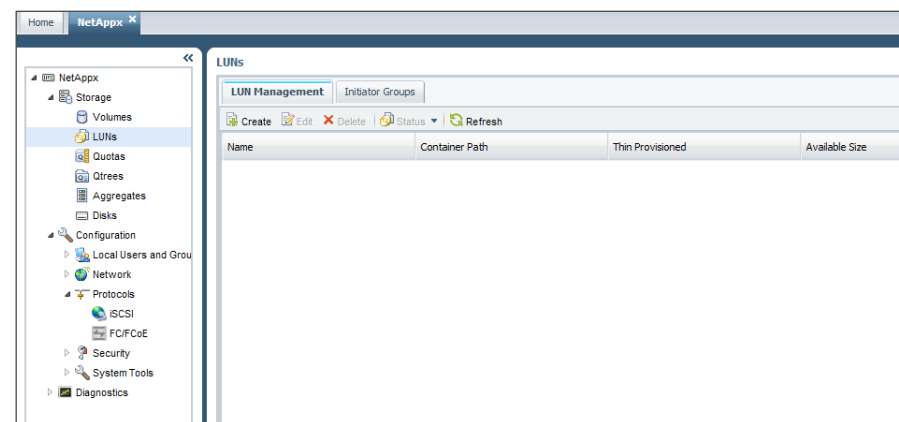
Before you can view or configure FC/FCoE, or iSCSI storage, you must add the licenses.

Step 1: In System Manager, navigate to **Configuration > System Tools > Licenses**, and then in the right pane, click **Add**.

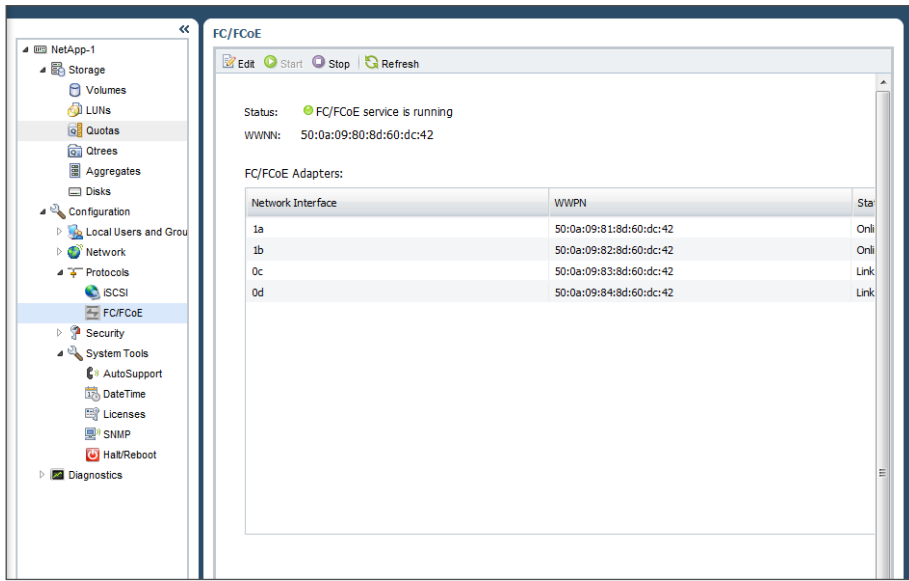
Step 2: On the Add License screen, in the **New license key** box, enter the license keys, click **Add**, and then **Click OK for the message that pops up for Add License**. If you are using iSCSI and FCoE, you may have multiple license keys.



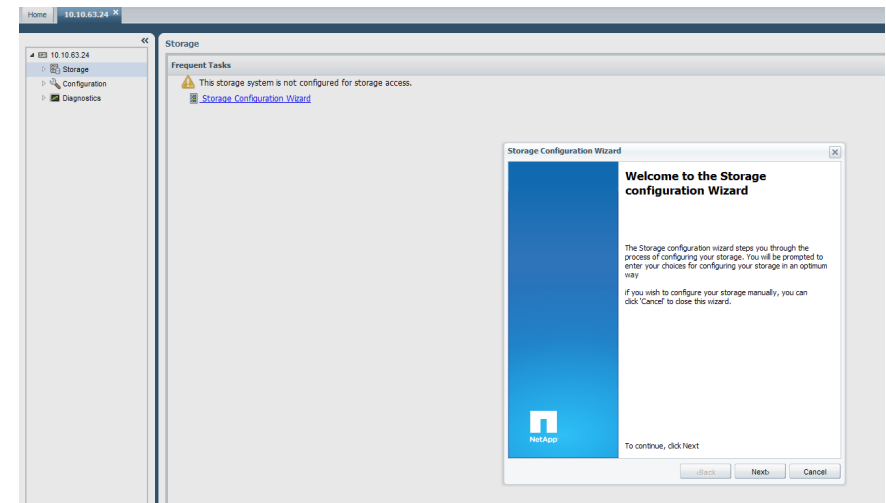
Step 3: After adding licenses in System Manager ensure that under **Storage**, LUNs appear, and under **Configuration > Protocols**, the necessary protocols appear.



Step 4: Navigate to **Configuration > Protocols > FC/FCoE**, and then in the right pane, click **Start**. If the FCoE service on the storage system is already running, the Start button is grayed-out. Alternatively, select and enable iSCSI if configuring iSCSI storage.



Step 1: In the left pane, select **Storage**, and in the right pane, click **Storage Configuration Wizard**, and then click **Next**.

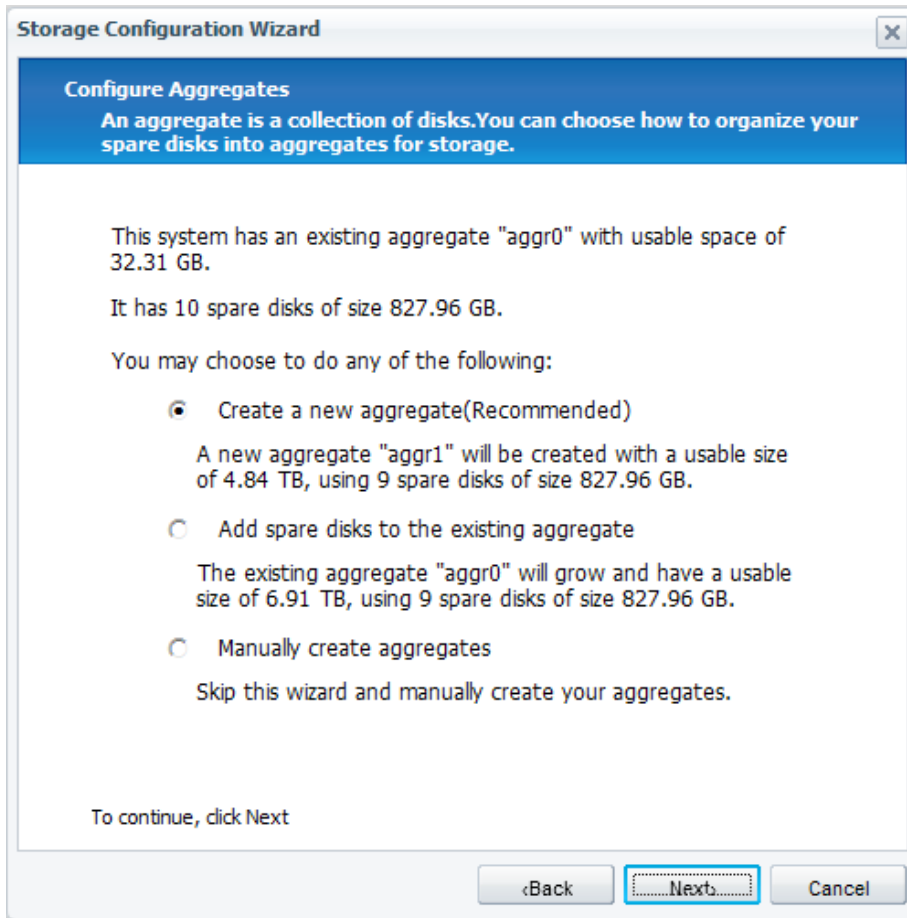


Procedure 2 Create an aggregate

An aggregate is NetApp's virtualization layer, which abstracts physical disks on a storage device from logical data sets that are referred to as flexible volumes. Aggregates are the means by which the total I/O operations per second (IOPS) available from all of the individual physical disks are pooled as a resource. Aggregates are well-suited to meet users' differing security, backup, performance, and data-sharing needs, as well as the most unpredictable and mixed workloads.

NetApp recommends that, whenever possible, you use a separate, small aggregate with RAID-DP to host the root volume. This stores the files required for the GUI management tools for the NetApp storage system. Place the remaining storage into a small number of large aggregates. This provides optimal performance because of the ability of a large number of physical spindles to service I/O requests. On smaller arrays, it may not be practical to have more than a single aggregate, due to the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

Step 2: On the Storage Configuration Wizard page, select **Create a new aggregate (Recommended)**, and then click **Next**.



The screenshot shows the 'Storage Configuration Wizard' window with the 'Configure Aggregates' tab selected. The header bar is blue with white text. Below the header, there is a description of aggregates. The main content area contains text about the existing aggregate 'aggr0' and three radio button options. The 'Next' button is highlighted with a blue border.

Storage Configuration Wizard

Configure Aggregates
An aggregate is a collection of disks. You can choose how to organize your spare disks into aggregates for storage.

This system has an existing aggregate "aggr0" with usable space of 32.31 GB.
It has 10 spare disks of size 827.96 GB.

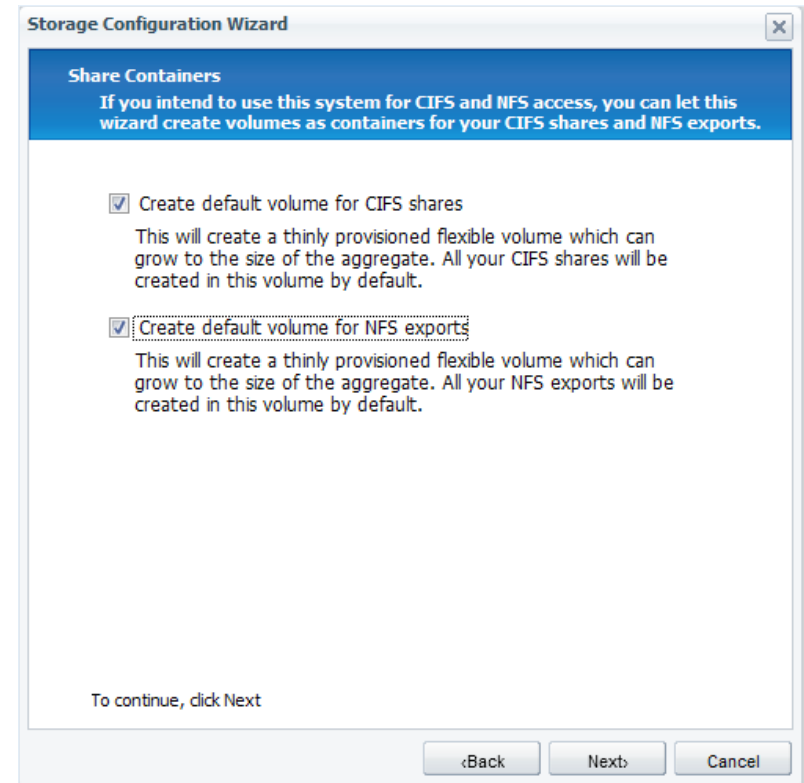
You may choose to do any of the following:

- ☒ **Create a new aggregate(Recommended)**
A new aggregate "aggr1" will be created with a usable size of 4.84 TB, using 9 spare disks of size 827.96 GB.
- ☐ Add spare disks to the existing aggregate
The existing aggregate "aggr0" will grow and have a usable size of 6.91 TB, using 9 spare disks of size 827.96 GB.
- ☐ Manually create aggregates
Skip this wizard and manually create your aggregates.

To continue, click Next

«Back **Next** Cancel

Step 3: If your configuration requires that you create default CIFS shares and NFS exports, select the check boxes, and then click **Next**.



The screenshot shows the 'Storage Configuration Wizard' window with the 'Share Containers' tab selected. The header bar is blue with white text. Below the header, there is a description of share containers and two checked checkboxes. The 'Next' button is highlighted with a blue border.

Storage Configuration Wizard

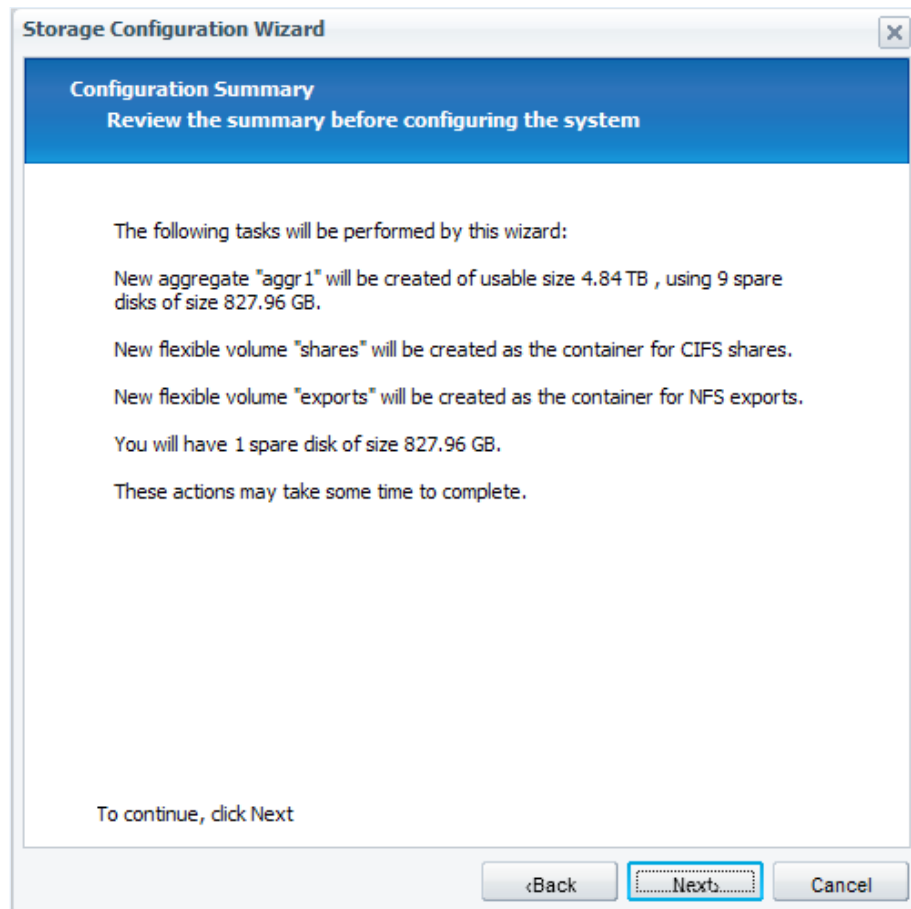
Share Containers
If you intend to use this system for CIFS and NFS access, you can let this wizard create volumes as containers for your CIFS shares and NFS exports.

- ☒ **Create default volume for CIFS shares**
This will create a thinly provisioned flexible volume which can grow to the size of the aggregate. All your CIFS shares will be created in this volume by default.
- ☒ **Create default volume for NFS exports**
This will create a thinly provisioned flexible volume which can grow to the size of the aggregate. All your NFS exports will be created in this volume by default.

To continue, click Next

«Back **Next** Cancel

Step 4: On the Configuration Summary page, verify the settings, and then click **Next**.



Step 5: Click **Finish**. A window appears with a message that the creation of an aggregate takes place in the background, allowing you to begin storing data immediately.

Step 6: Click **OK**. The right pane of the Storage screen is refreshed with Create Aggregate, Create Volume, and Create Qtree links.



Reader Tip

For instructions about how to manually configure an aggregate, see the Creating an Aggregate section in the *Data ONTAP Storage Management Guide* here: <http://hd.kvsconsulting.us/netappdoc/801docs/pdfs/ontap/smg.pdf>

Procedure 3

Configure flexible volumes

FlexVol volumes are thin storage containers that can contain LUNs and/or file shares that are accessed by servers over FC, FCoE, iSCSI, NFS, or CIFS. A FlexVol volume is a virtual volume that you can manage and move independently from physical storage. It can be created and resized larger or smaller as your application needs change.

Step 1: Navigate to **Storage > Volumes**, and then click **Create**.

Step 2: On the General tab, configure the following properties, and then click **Create**.

- Name—Enter a name for the volume. NetApp recommends that you use a combination of the hostname, physical disk, and replication policy. (Example: HostB_Disk1_4hmirror)
- Aggregate— Specify an aggregate for the LUN. Keep the value **aggr1** that gets populated.
- Storage Type—**SAN**
- Total Size—Enter the capacity size of the LUN.
- Snapshot reserve—Select the percentage of space in the LUN that you want to dedicate for Snapshot data. The default setting is 5%, which means that 5% of the LUN will not be available for other data storage. You can choose a higher value or as little as 0%, depending on your requirements. NetApp recommends configuring all volumes with 0% and disabling the default Snapshot schedule.
- Thin Provisioned—Space reclamation. Select this option if you want to create a thinly provisioned volume that consumes space on the disk only as data is written.

Step 3: Click Create.

Create Volume

General

Name: vol_NetAppx_Disk21

Aggregate: aggr1 Choose

Free space in aggregate is 4.84 TB

Storage Type

☒ SAN (Used for FC/FCoE or iSCSI access)

Size

Total Size: 20 GB

Snapshot Reserve (%): 0

Data Space: 20 GB

Snapshot Space: 0 Byte

Thin Provisioning

☒ Thin Provisioned

Allocate space as it's used, otherwise, allocate the space now. Recommended for increasing utilization when each volume is unlikely to use all of its allocated space.

[Tell me more about Thin Provisioning](#)

Create Cancel



Tech Tip

If you are using this volume for Cisco Unified Communications services, such as Cisco Unified Communications Manager or Cisco Unity Connection, Cisco recommends that you choose thick provisioning. For a description of thin provisioning, see the Increasing Efficiency and Flexibility with Advanced Features section of this guide.

Process

Deploying FCoE or iSCSI Storage

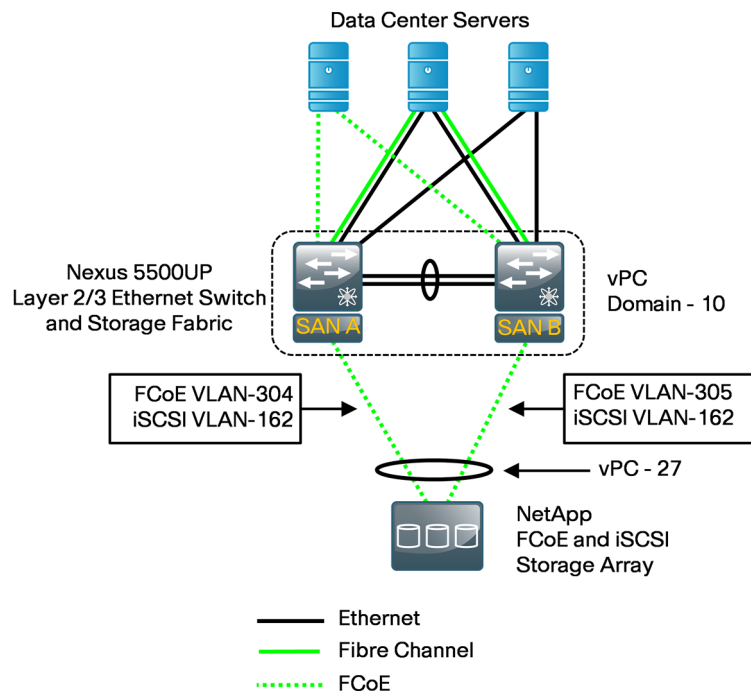
1. Configure the data-center core
2. Add an initiator group

The data-center core deployed in the *Cisco SBA—Data Center Deployment Guide* provides an Ethernet transport that supports IP traffic, as well as FCoE. The Cisco SBA data-center core can also support native FC. This guide uses the Ethernet transport to connect the NetApp storage array to the network for iSCSI or FCoE operation.

Procedure 1

Configure the data-center core

Before you configure the storage controller for connectivity to the servers, you must prepare the data-center core network. Details for configuring Ethernet and FCoE in the Cisco SBA data-center core are provided in the *Cisco SBA—Data Center Deployment Guide*. This procedure shows you how to prepare for FCoE and iSCSI storage, and assumes that the Cisco Nexus data-center core switches have been configured for virtual port channel (vPC) operations according to the *Cisco SBA—Data Center Deployment Guide*.



Follow the steps in this procedure to configure both FCoE and iSCSI traffic. If you are configuring for iSCSI traffic only, perform Step 7 through Step 10, and then skip to Procedure 2, “Add an initiator group.”

Step 1: Log in to the first Cisco Nexus 5500UP data center core switch.

Step 2: Ensure that you have licenses available for the FCoE ports to be used. An FCoE license is required to keep FCoE enabled past the 90-day grace period.

```
dc5548ax# install license bootflash://license.lic
```

Step 3: Verify that FCoE, LACP, and vPC are enabled.

```
dc5548ax# show feature
```

Feature Name	Instance	State
fcoe	1	enabled
lACP	1	enabled
vpc	1	enabled

Step 4: If a feature in Step 3 is not enabled, enable it.

Example

```
feature fcoe
```

Step 5: If you have not already configured your Cisco Nexus 5500UP switches for QoS by following the Configuring the Data Center Core procedure in the *Cisco SBA—Data Center Deployment Guide*, you must enable quality of service (QoS) for FCoE operation on the Cisco Nexus 5500UP.

Four lines of QoS statements map the baseline system QoS policies for FCoE. Without these commands, the virtual FC interface will not function when activated. If you followed the *Cisco SBA—Data Center Deployment Guide to deploy your network*, you should have already executed a more comprehensive QoS policy, which includes FCoE traffic classification, so you can skip this step. If you use the commands below for the baseline FCoE QoS operation, you will overwrite your existing QoS policy.

```
system qos
service-policy type qos input fcoe-default-in-policy
service-policy type queuing input fcoe-default-in-policy
service-policy type queuing output fcoe-default-out-policy
service-policy type network-qos fcoe-default-nq-policy
end
```



Tech Tip

All FC and FCoE control and data traffic is automatically classified into the FCoE system class, which provides a no-drop service. On the Cisco Nexus 5010 and Cisco Nexus 5020, this class is created automatically when the system starts up. The class is named class-fcoe in the CLI.

Step 6: Create a VLAN that carries FCoE traffic to the storage controller, and map it to the VSAN for the respective SAN fabrics.

Table 1 - FC VSAN to FCoE VLAN mapping

Data center core switch	VSAN	FCoE VLAN
Nexus 5548UP-1	4	304
Nexus 5548UP-2	5	305

- On the first Cisco Nexus 5500UP, map VLAN 304 to VSAN 4. VLAN 304 carries all VSAN 4 traffic over the trunk.

```
vsan database
vsan 4 name General-Storage
vlan 304
fcoe vsan 4
exit
```

- On the second Cisco Nexus 5500UP, map VLAN 305 to VSAN 5.

```
vsan database
vsan 5 name General-Storage
vlan 305
fcoe vsan 5
exit
```

Step 7: Create a VLAN on each Cisco Nexus 5500UP switch to carry iSCSI traffic.

```
vlan 162
name iSCSI
```

Step 8: Create a Layer 3 switch virtual interface (SVI) on each Cisco Nexus 5500UP switch to provide a default route for the iSCSI IP subnet.

- Configure the first Cisco Nexus 5500UP switch.

```
interface Vlan162
no shutdown
no ip redirects
ip address 10.4.62.2/24
hsrp 162
priority 110
ip 10.4.62.1
```

- Configure the second Cisco Nexus 5500UP switch.

```
interface Vlan162
no shutdown
no ip redirects
ip address 10.4.62.3/24
hsrp 162
ip 10.4.62.1
```

Step 9: Configure a vPC port channel to connect to the storage controller to transport both iSCSI and FCoE traffic. If you are deploying iSCSI only, you can omit VLAN 304 or 305 on the trunk.

- Configure the first Cisco Nexus 5500UP switch.

```
interface port-channel 27
switchport mode trunk
switchport trunk native vlan 162
switchport trunk allowed vlan 162,304
spanning-tree port type edge trunk
vpc 27
```

- Configure the second Cisco Nexus 5500UP switch.

```
interface port-channel 27
switchport mode trunk
switchport trunk native vlan 162
switchport trunk allowed vlan 162,305
spanning-tree port type edge trunk
vpc 27
```

Step 10: Map the physical port connected to the storage array on each Cisco Nexus 5500UP switch to the port channel. If you are deploying iSCSI only, you can omit VLAN 304 or 305 on the trunk.

- Configure the first Cisco Nexus 5500UP switch.

```
interface ethernet1/27
channel-group 27 mode active
switchport mode trunk
switchport trunk native vlan 162
switchport trunk allowed vlan 162,304
```

- Configure the second Cisco Nexus 5500UP switch.

```
interface ethernet1/27
channel-group 27 mode active
switchport mode trunk
switchport trunk native vlan 162
switchport trunk allowed vlan 162,305
```

If you are deploying iSCSI only, you can proceed to Procedure 2, “Add an initiator group.”

Step 11: Create a virtual fibre-channel (vFC). The vFC transports FCoE traffic to the storage controller.

- Configure the first Cisco Nexus 5500UP switch.

```
interface vfc27
bind interface port-channel 27
switchport trunk allowed vsan 4
no shutdown
vsan database
vsan 4 interface vfc 27
```

- Configure the second Cisco Nexus 5500UP switch.

```
interface vfc27
bind interface port-channel 27
switchport trunk allowed vsan 5
no shutdown
vsan database
vsan 5 interface vfc 27
```



Tech Tip

A vFC port can be bound to a specific Ethernet port; this is done when a vPC has not been deployed in the design. A vFC port can be bound to a specific port world-wide name (PWWN) when an FIP snooping bridge is attached. In the case of a vPC, you are binding to port channel as a requirement for vPC and LACP to work with the converged network adapter. You do not have to assign the same number to the vFC as the physical port number, but doing so improves manageability.

Step 12: Obtain the FC PWWN for the storage controller (target) and the server (initiator) to configure zoning. To obtain the storage controller PWWNs, enter the following command on the storage controller CLI.

```
NetAppx> fcp portname show
Portname                      Adapter
-----
50:0a:09:81:8d:60:dc:42      1a
50:0a:09:82:8d:60:dc:42      1b
```

You can also use the **show flogi database** and **show fcns database** commands on the Cisco Nexus 5500UP switches to display active PWWNs on the FC fabric. The FC initiators and targets should be enabled and connected to the FC network to display the active nodes.

Step 13: Create the FC zoneset and zone for FCoE operation.

- Configure the first Cisco Nexus 5500UP switch.

```
zone name p30-ucsc200m2-2-vhba3_netapp-e1a vsan 4
member pwn 50:0a:09:81:8d:60:dc:42
member pwn 20:00:cc:ef:48:ce:c1:f9
exit
```

```
zoneset name FCOE_4 vsan 4
member p30-ucsc200m2-2-vhba3_netapp-e1a
exit
```

```
zoneset activate name FCOE_4 vsan 4
```

- Configure the second Cisco Nexus 5500UP switch.

```
zone name p30-ucsc200m2-2-vhba4_netapp-e1b vsan 5
member pwn 50:0a:09:82:8d:60:dc:42
member pwn 20:00:cc:ef:48:ce:c1:fa
exit
```

```
zoneset name FCOE_5 vsan 5
member p30-ucsc200m2-2-vhba4_netapp-e1b
exit
```

```
zoneset activate name FCOE_5 vsan 5
```

Step 14: Display active zoneset members. The asterisk (*) at the beginning of the **fcid** output line indicates an active member. Note that you will have to complete the NetApp FCoE configuration section before the active member is logged in to the FC fabric and becomes active.

- Configure the first Cisco Nexus 5500UP switch.

```
dc5548a# show zoneset active vsan 4
zoneset name FCOE_4 vsan 4
  zone name p30-ucsc200m2-2-vhba3_netapp-e1a vsan 4
    * fcid 0x13000e [pwwn 50:0a:09:81:8d:60:dc:42]
    * fcid 0x130006 [pwwn 20:00:cc:ef:48:ce:c1:f9]
```

- Configure the second Cisco Nexus 5500UP switch.

```
dc5548b# show zoneset active vsan 5
zoneset name FCOE_5 vsan 5
  zone name p30-ucsc200m2-2-vhba4_netapp-e1b vsan 5
    * fcid 0x93000e [pwwn 50:0a:09:82:8d:60:dc:42]
    * fcid 0x930006 [pwwn 20:00:cc:ef:48:ce:c1:fa]
```

Procedure 2 Add an initiator group

Initiator groups, which are commonly referred to as igroups, access control lists, or LUN masking lists, are a way to group LUNs according to the parameters of the environments or servers they service. You create an initiator group by specifying a collection of PWWNs of initiators in an FC network. You use the PWWNs of the host's HBAs. Initiator groups can be created during LUN creation or independently, but when you first create a LUN, it is helpful if the appropriate initiator group already exists so that it can be associated to the LUN directly within the LUN creation wizard. Initiators can be programmed for FCoE or iSCSI operation. An iSCSI or FCoE initiator group may be associated to a LUN for iSCSI or FCoE access, and holds one or more initiator IDs.

Before you can view LUNs or initiator groups on the NetApp filer, you must have added licenses for FC/FCoE and iSCSI by following Procedure 1, in the Provisioning Storage section of this guide.

Step 1: Navigate to **Storage > LUNs**, and in the right pane, click the **Initiator Groups** tab, and then click **Create**.

Step 2: Click the **General** tab, and then make the following changes:

- Name—Enter a name for the initiator group
- Operating System—Choose the appropriate operating system for the host that will be accessing this storage LUN
- Type—Select the type of protocol to be used
- Enable ALUA—Leave this check box cleared

Create Initiator Group

General | Initiators

Name:

Operating System:

Type

Select the supported protocol for this group

☐ iSCSI

☒ FC/FCoE

ALUA

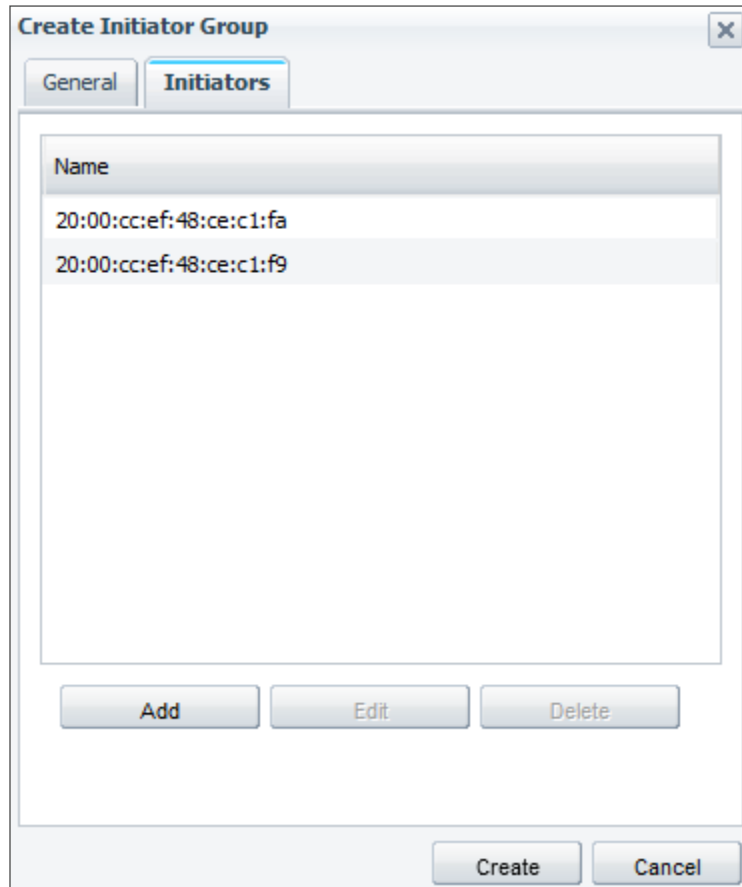
ALUA helps in managing multipath IOs

☐ Enable ALUA (Asymmetric Logical Unit Access)

Create Cancel

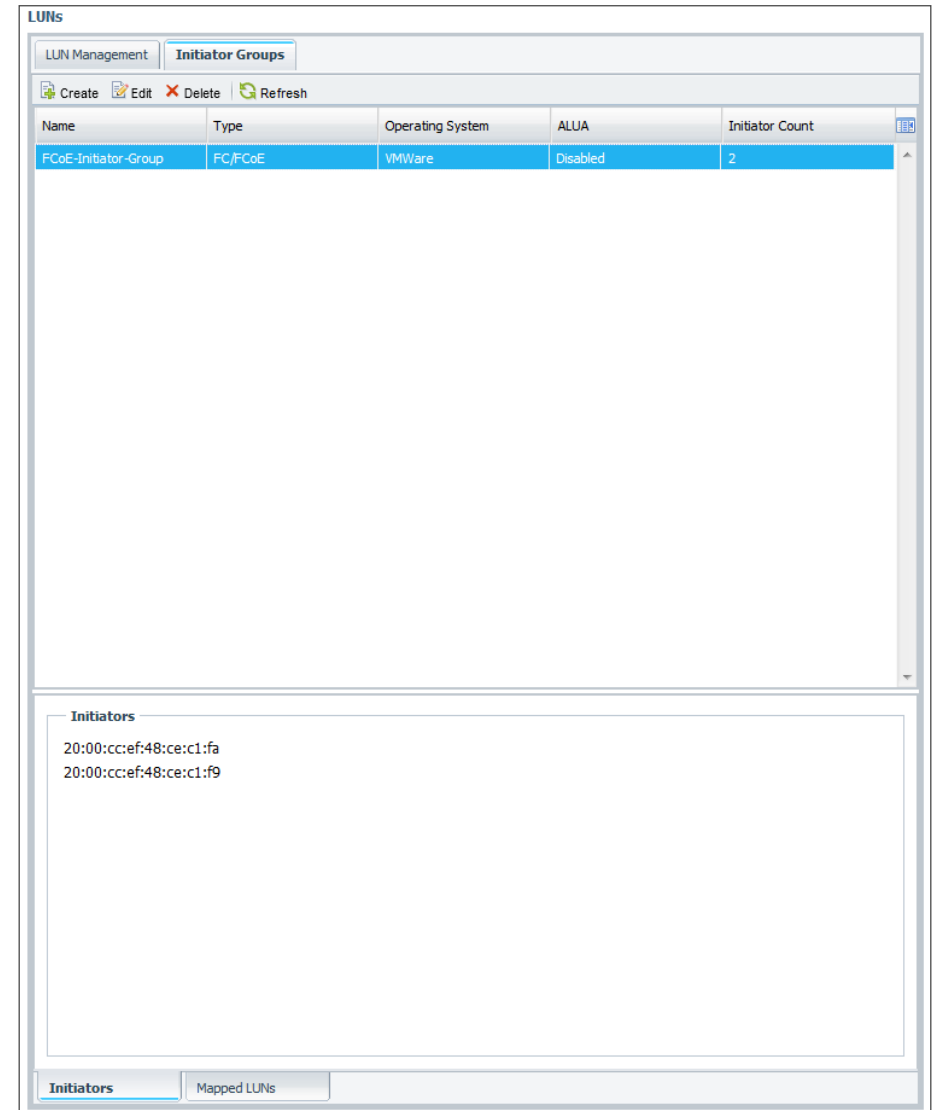
Step 3: Click the **Initiators** tab, and then click **Add**.

- If you are creating an FCoE initiator group, the name of the initiator must match the PWWN identified in the Cisco Nexus 5500UP zone configuration of the host machine that is accessing storage. You configured this in Step 13 of Procedure 1, "Configure the data-center core."
- If you are creating an iSCSI initiator group, add the initiator's world-wide name (WWN) or iSCSI qualified node (IQN) name for iSCSI initiator group. The iSCSI initiator must be enabled, active, and pointed to the NetApp filer (target).



The "Create Initiator Group" dialog box is shown with the "Initiators" tab selected. It contains a list of WWN addresses: 20:00:cc:ef:48:ce:c1:fa and 20:00:cc:ef:48:ce:c1:f9. Below the list are "Add", "Edit", and "Delete" buttons. At the bottom are "Create" and "Cancel" buttons.

Step 4: Click **Create**. After the group is created, the group name in the upper half of the right pane is highlighted, and initiators that were added to the group are shown in the bottom half of the screen.



The "LUN Management" window shows the "Initiator Groups" tab. A table lists the created group:

Name	Type	Operating System	ALUA	Initiator Count
FCoE-Initiator-Group	FC/FCoE	VMWare	Disabled	2

Below the table, the "Initiators" section lists the WWN addresses: 20:00:cc:ef:48:ce:c1:fa and 20:00:cc:ef:48:ce:c1:f9. The "Mapped LUNs" tab is also visible at the bottom.



Tech Tip

To determine the target WWN from the controller using a console CLI, use the **fc** or **iscsi nodename** command.

```
NetAppx> iscsi nodename
iSCSI target nodename: iqn.1992-08.com.
netapp:sn.151741374
NetAppx> fcp portname show
Portname                Adapter
-----                -
50:0a:09:81:8d:60:dc:42  1a
50:0a:09:82:8d:60:dc:42  1b
```

Process

Adding a LUN

1. Create a LUN

You can create the following four types of storage objects on the same NetApp storage system:

- Aggregate
- Volume
- LUN
- Qtree

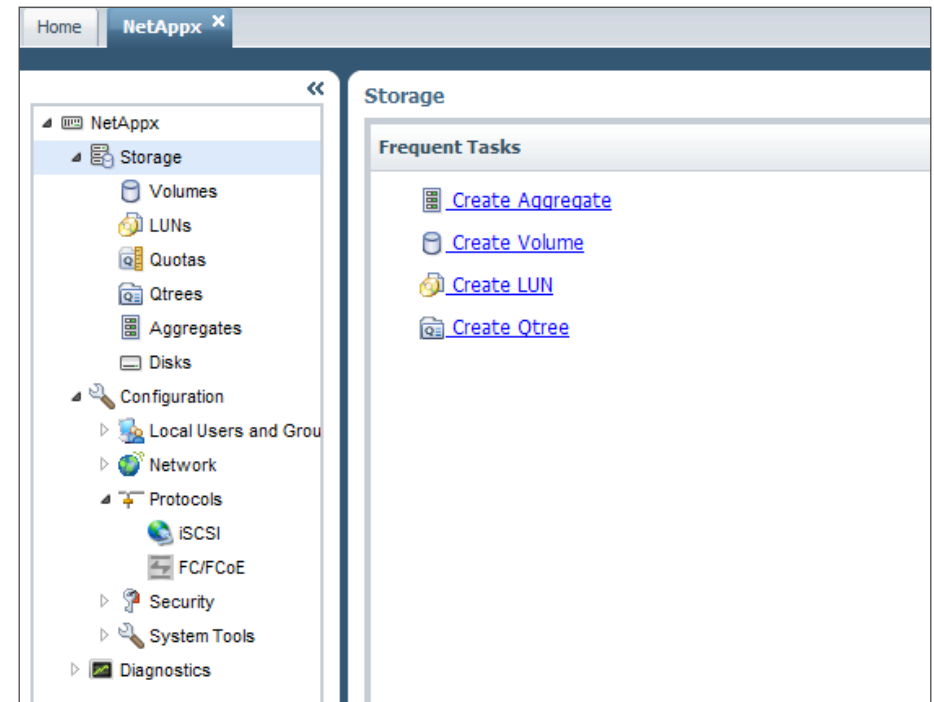
Creating a LUN automatically creates a default volume as part of the process. To configure a volume with properties other than the defaults, see Procedure 3, "Configure flexible volumes."

Procedure 1

Create a LUN

LUNs are logical units of storage provisioned from a NetApp storage system directly to servers. Hosts can access the LUNs as physical disks using FC, FCoE, or iSCSI protocols. The following steps illustrate how to configure an FCoE LUN.

Step 1: In the left pane, select **Storage**, and then in the right pane, click **Create LUN**.



Step 2: On the Create LUN Wizard Welcome page, click **Next**.

Step 3: On the General Properties page, enter the following configuration details, and then click **Next**.

- Name—LUN name
- Description (optional)—Description for the LUN
- Type—Operating system type
- Size—LUN size
- Thin Provisioned—Select this check box unless your deployment meets the exception discussed in the following Tech Tip



Tech Tip

If you are using this volume for Cisco unified communications services, such as Cisco Unified Communications Manager or Cisco Unity Connection, Cisco recommends that you choose thick provisioning. For a description of thin provisioning, see the “Increasing Efficiency and Flexibility with Advanced Features” section of this guide.

Each operating system (OS) maps data to LUNs slightly differently. The OS-type parameter determines the on-disk layout of the LUN. It is important to specify the correct OS-type to make sure that the LUN is properly aligned with the file system on it. This is because optimal performance with the storage system requires that I/O is aligned to a 4096-byte boundary. If an I/O is unaligned, it can:

- Cause an increase in per-operation latency.
- Require the storage system to read-from or write-to more blocks than necessary to perform logical I/O.

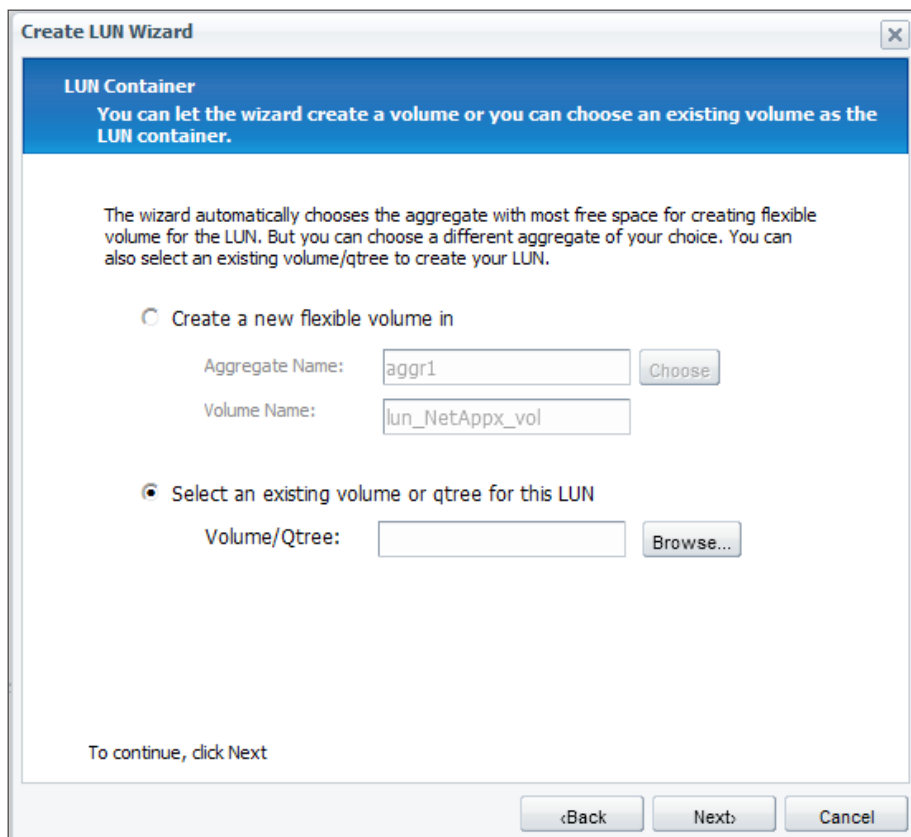
This issue is not unique to NetApp storage. Any storage vendor or host platform can exhibit this problem. After the LUN is created, you cannot modify the LUN host OS type.



Tech Tip

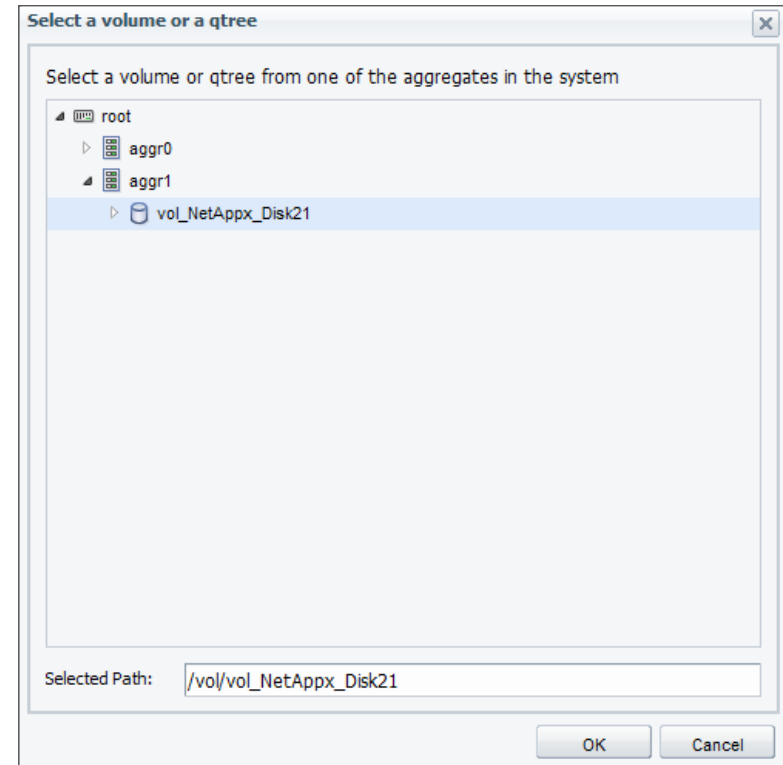
If you have selected the incorrect OS type and then created the LUN, the LUN at the storage array will be misaligned. In order to correct this problem, you must create a new LUN, and then select the correct OS type.

Step 4: Select the **Select an existing volume or qtree for this LUN** option, and then click **Browse**.



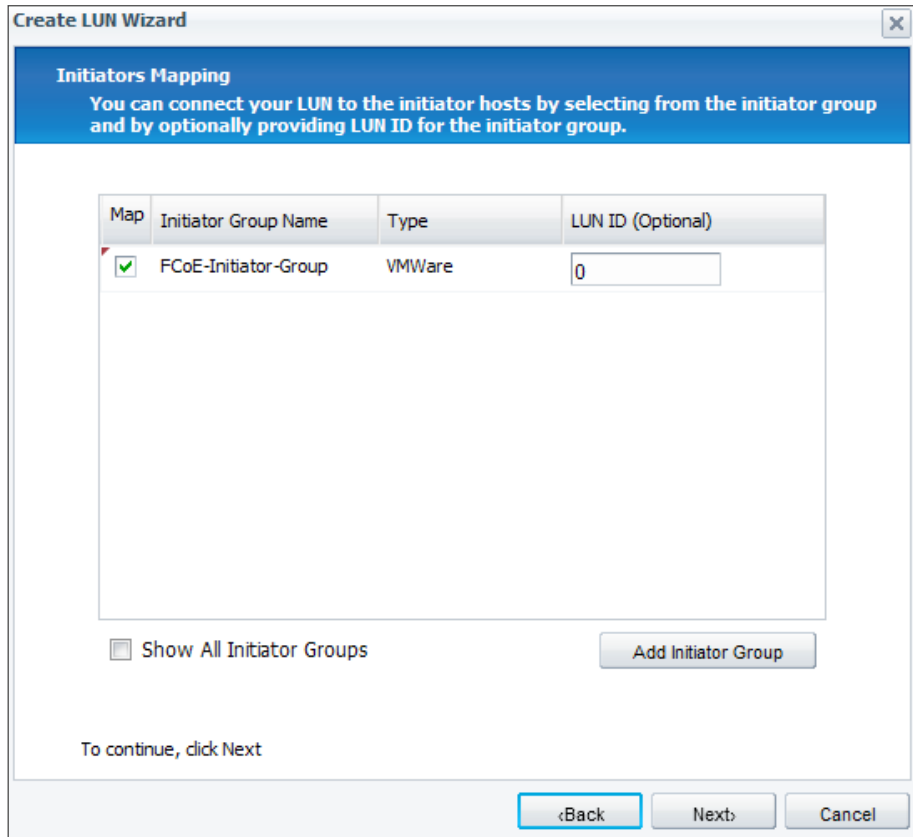
The **Create LUN Wizard** dialog box has a title bar with a close button. Below the title bar is a blue header section with the text "LUN Container" and "You can let the wizard create a volume or you can choose an existing volume as the LUN container." Below this is a paragraph of text: "The wizard automatically chooses the aggregate with most free space for creating flexible volume for the LUN. But you can choose a different aggregate of your choice. You can also select an existing volume/qtree to create your LUN." There are two radio button options. The first is "Create a new flexible volume in" with input fields for "Aggregate Name:" (containing "aggr1") and "Volume Name:" (containing "lun_NetAppx_vol"), and a "Choose..." button. The second option is "Select an existing volume or qtree for this LUN", which is selected. It has a "Volume/Qtree:" input field and a "Browse..." button. At the bottom, there is a "To continue, click Next" instruction and three buttons: "Back", "Next", and "Cancel".

Step 5: In the **Select a volume or a qtree** window, drill down to the location of the volume that you configured in Step 2 of Procedure 3 "Configure flexible volumes," click **OK**, and then click **Next**.



The **Select a volume or a qtree** dialog box has a title bar with a close button. Below the title bar is the text "Select a volume or qtree from one of the aggregates in the system". It features a tree view showing a hierarchy: "root" (expanded) contains "aggr0" and "aggr1" (expanded). "aggr1" contains "vol_NetAppx_Disk21", which is selected and highlighted in blue. Below the tree view is a "Selected Path:" text field containing the path "/vol/vol_NetAppx_Disk21". At the bottom right are "OK" and "Cancel" buttons.

Step 6: By default this shows the initiator group (host) that was created earlier. Click **Create** to open the Create Initiator Group window. Under Map, check the box, and then, if you are creating an operating system boot LUN, enter **0** in the **LUN ID** box. If you are creating any other type of LUN, enter an unused LUN ID or allow Data ONTAP to assign the LUN ID automatically.



The screenshot shows the 'Create LUN Wizard' window, specifically the 'Initiators Mapping' step. The window has a blue header with the title 'Initiators Mapping' and a subtitle: 'You can connect your LUN to the initiator hosts by selecting from the initiator group and by optionally providing LUN ID for the initiator group.' Below the header is a table with four columns: 'Map', 'Initiator Group Name', 'Type', and 'LUN ID (Optional)'. The first row shows a checked box in the 'Map' column, 'FCoE-Initiator-Group' in the 'Initiator Group Name' column, 'VMWare' in the 'Type' column, and '0' in the 'LUN ID (Optional)' column. Below the table is a checkbox labeled 'Show All Initiator Groups' and a button labeled 'Add Initiator Group'. At the bottom of the window, there is a button labeled 'Back' (highlighted with a red box), a button labeled 'Next', and a button labeled 'Cancel'. A message at the bottom left says 'To continue, click Next'.

Map	Initiator Group Name	Type	LUN ID (Optional)
<input checked="" type="checkbox"/>	FCoE-Initiator-Group	VMWare	0

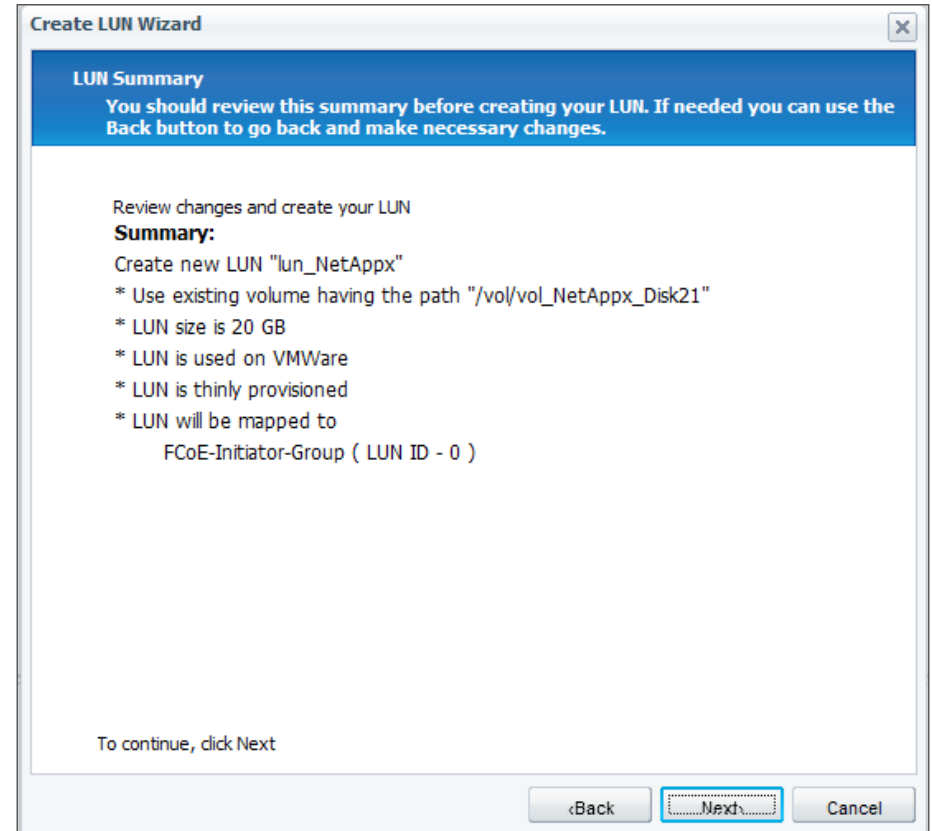
☐ Show All Initiator Groups Add Initiator Group

To continue, click Next

Back Next Cancel

Step 7: Click **Next**.

Step 8: On the LUN Summary page, verify the settings, and then click **Next**.



The screenshot shows the 'Create LUN Wizard' window, specifically the 'LUN Summary' step. The window has a blue header with the title 'LUN Summary' and a subtitle: 'You should review this summary before creating your LUN. If needed you can use the Back button to go back and make necessary changes.' Below the header, the text reads: 'Review changes and create your LUN'. Under the heading 'Summary:', it says 'Create new LUN "lun_NetAppx"'. Below this, there is a list of bullet points: '* Use existing volume having the path "/vol/vol_NetAppx_Disk21"', '* LUN size is 20 GB', '* LUN is used on VMWare', '* LUN is thinly provisioned', and '* LUN will be mapped to FCoE-Initiator-Group (LUN ID - 0)'. At the bottom of the window, there is a message 'To continue, click Next' and three buttons: 'Back', 'Next' (highlighted with a red box), and 'Cancel'.

Review changes and create your LUN

Summary:

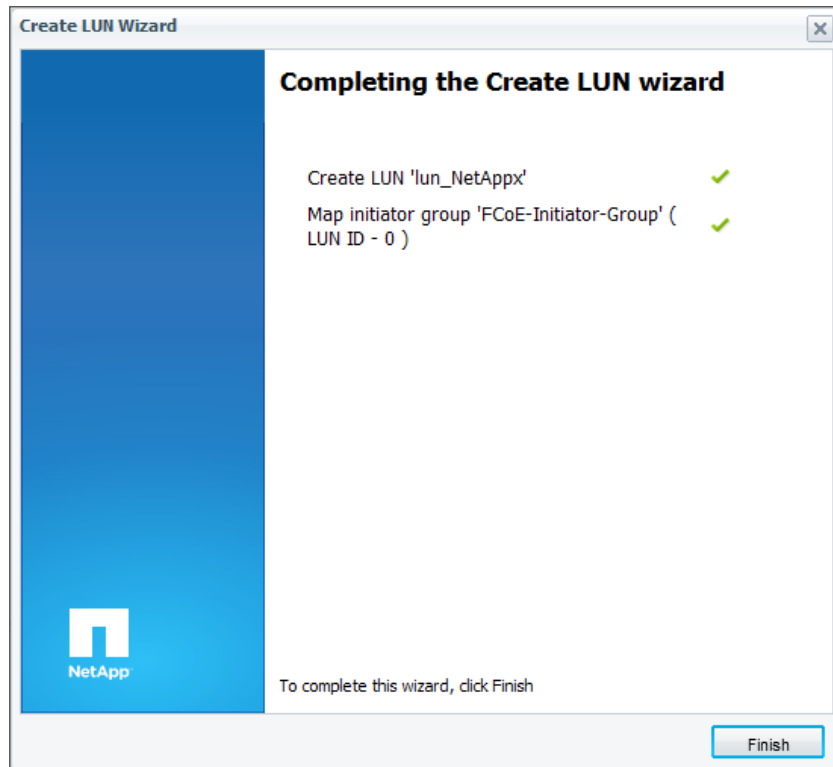
Create new LUN "lun_NetAppx"

- * Use existing volume having the path "/vol/vol_NetAppx_Disk21"
- * LUN size is 20 GB
- * LUN is used on VMWare
- * LUN is thinly provisioned
- * LUN will be mapped to FCoE-Initiator-Group (LUN ID - 0)

To continue, click Next

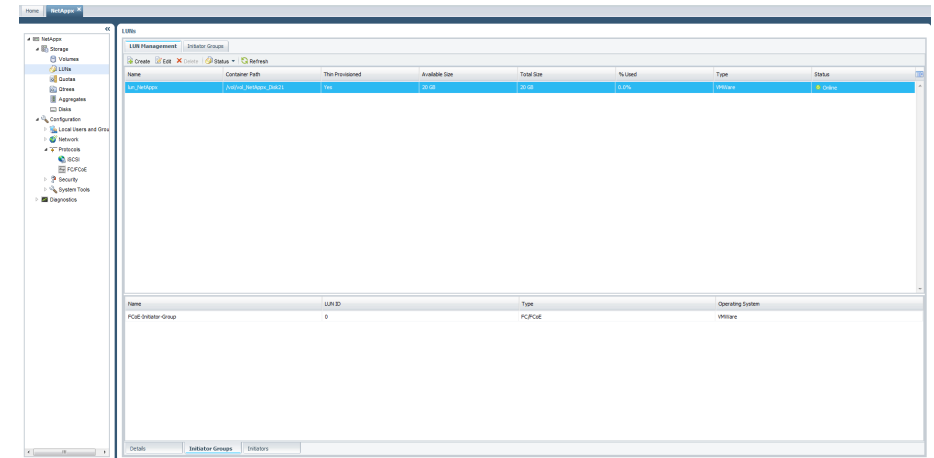
Back Next Cancel

Step 9: On the Completing the Create LUN wizard page, click **Finish**.



After the LUN has been created, it is accessible to the FCoE host that was identified in the initiator mapping.

Step 10: To view the LUN that you created, navigate to **Storage > LUNs**, and in the right pane, click the **LUN Management** tab. The LUN is listed with its Status as Online in the top half of the screen. You can view the LUN's **Initiator Groups** and **Initiators** by clicking the respective tabs in the bottom half of the screen.



Step 11: This concludes the NetApp configuration for FCoE or iSCSI.

Step 12: You now have an end-to-end storage configuration in place. You can activate any of the available Ethernet protocols (NFS, CIFS, iSCSI) and you can access them from the VIF IP addresses. In addition, you can map to any iSCSI LUNs created on the controller(s) that you have zoned to.

Increasing Efficiency and Flexibility with Advanced Features

Thin Provisioning

Traditional storage provisioning and preallocation of storage on disk are methods that storage administrators understand well. It is a common practice for server administrators to overprovision storage to avoid running out of storage, and thereby avoid the associated application downtime when they expand the provisioned storage to new levels.

Although no system can run at 100% storage utilization, there are storage virtualization methods that allow administrators to address and oversubscribe storage in the same manner as server resources (such as CPU, memory, and networking). This form of storage virtualization is referred to as thin provisioning.

While traditional provisioning preallocates storage, thin provisioning provides storage on-demand. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual application requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. The drawback to thin provisioning and oversubscribing storage is that, without the addition of physical storage, if every application requires its maximum possible storage at the same time, there will not be enough storage to satisfy the requests.

NetApp FlexVol uses thin provisioning to allow LUNs that are presented as physical disks to be provisioned to their total capacity, yet consume only as much physical storage as is required to store data. LUNs connected as pass-through disks can also be thin provisioned. Thin provisioning applies equally to file shares.

NetApp Deduplication and Compression

With NetApp deduplication, server deployments can eliminate the duplicate data in their environment, enabling greater storage utilization. Deduplication can be seamlessly introduced into the server environment without having to make any changes to server administration, practices, or tasks. Deduplication runs on the NetApp storage system at scheduled intervals and does not consume any CPU cycles on the server.

Deduplication can be extremely helpful for virtual server scenarios such as fixed-size virtual hard drives, frequent creation and deletion of virtual disk files on the SAN LUNs, and data in the child VM.

Deduplication is enabled on the NetApp volume, and the amount of data deduplication realized is based on the commonality of the data stored in a deduplication-enabled volume.

NetApp data compression is a new feature that compresses data as it is written to NetApp FAS and V-Series storage systems. Like deduplication, NetApp data compression works in both SAN and NAS environments, and is application and storage tier agnostic.

NetApp Snapshot

A NetApp Snapshot copy is a locally retained, frozen, space-efficient, read-only view of a volume or an aggregate. Its improved stability, scalability, recoverability, and performance make it more efficient than other storage snapshot technologies.

Snapshot copies facilitate frequent low-impact, user-recoverable online backup of files, directory hierarchies, LUNs, and application data. They offer a secure and simple method of restoring data so that users can directly access the Snapshot copies and recover from accidental file deletion, data corruption, or modification. The SnapManager suite of products, which is available for various enterprise applications, uses the features of Snapshot copies and delivers an enterprise-class data protection solution.

NetApp FlexClone

NetApp FlexClone technology creates true cloned volumes, which are instantly replicated data sets, files, LUNs, and volumes that use no additional storage space at the time of creation. A FlexClone volume is a writable point-in-time copy generated from the Snapshot copy of a FlexVol volume. It has all the features of a FlexVol volume, including growing, shrinking, and being the base for a Snapshot copy or even another FlexClone volume.

FlexClone volumes deployed in a virtualized environment offer significant savings in dollars, space, and energy. Additionally, the performance of a FlexClone volume or file is identical to the performance of any other FlexVol volume or individual file.

Backup, Disaster Recovery, and High Availability

Backup and recovery are the most critical components of the data protection plan. A backup is crucial to protect and recover business information if data is changed unexpectedly, a system is compromised, or a site is lost.

NetApp backup and recovery solutions equip users to increase the reliability of data protection while minimizing management-overhead and cost. These solutions fit into any strategy, enabling users to meet their service-level requirements.

Backup and Recovery Concepts

Data protection plans for a virtualized environment become more critical as consolidation brings all of the crucial data into one place, so that any failure results in a massive impact on business applications.

Backup tasks running in the server virtualized infrastructure are often resource-intensive (CPU, memory, disk I/O, and network) and can result in bottlenecks that adversely affect the performance of the other business-critical applications that share the environment. Backup schedules must be closely coordinated with the applications that are running on the available resources.

Disaster Recovery

Business operations depend heavily on information systems and the related IT infrastructure. A minor application outage can significantly affect operations, and the effect of data loss is even more critical. There are various metrics that are commonly used in designing a business continuity plan. Two of the most frequently used metrics are recovery point objective (RPO) and recovery time objective (RTO). RPO, measured in minutes and hours, describes how far the recovered data are out of sync with the production data at the time of disaster. RTO, measured in minutes, describes how fast the operations can be restored.

Several approaches have been developed to increase data availability and business continuity in case of disaster occurring at the hardware or software level, and even site failures. Backup methods primarily provide a way to recover from data loss from an archived medium—a high-level data protection method.

Redundant hardware setups can provide second-level protection to mitigate damage caused by hardware failures. Data mirroring is another mechanism to increase data availability and minimize downtime.

NetApp offers the SnapMirror solution, which empowers IT infrastructures with a fast, flexible data replication mechanism over Ethernet and FC networks. It is a key component to consider when designing and deploying enterprise data protection plans. SnapMirror is an efficient data replication solution that takes advantage of underlying NetApp technologies, such as Snapshot, FlexClone, and deduplication. Disaster recovery is its primary objective, and SnapMirror can also assist in other critical application areas, such as disaster recovery testing, application testing, load sharing, remote tape archiving, and remote data access.

Business Continuity Concepts

Disaster can occur in any IT infrastructure, and a data protection plan is even more critical for environments that are consolidated by using server virtualization. This is true because consolidation adds complexity by sharing reduced physical hardware resources for the applications and the business-critical data that are running.

The infrastructure must be designed to pay special attention to challenges that can crop-up in a virtualized environment. These challenges include:

- Less time (or possibly no time) available to schedule downtime windows to perform cold-backup on virtual machines.
- Performing hot backup of virtual machines can result in inconsistent backup copies, which are of no use during recovery.
- Various operating system instances in the infrastructure, which can make it difficult to identify a consistent state for backup.
- Replicating data over LAN or WAN can consume twice as much of the available resources.
- Increased total cost of ownership (TCO) and unused infrastructure when planning for identical resources at the disaster recovery site.

NetApp Advanced Solutions

NetApp offers solutions that complement the server virtualization solutions and help to mitigate these challenges. Solutions such as NetApp Snapshot, FlexClone, compression, and deduplication enable an architect to design a complete data protection solution and to efficiently use available resources.

NetApp SnapMirror

NetApp SnapMirror software is a simple, flexible, cost-effective disaster-recovery and data-distribution solution that is deployed for more of the enterprise application infrastructure. Data is replicated across LAN or WAN, offering high availability and faster disaster-recovery for business-critical applications. Continuous data mirroring and mirror updates across multiple NetApp storage systems facilitate the mirrored data for multiple purposes. Businesses in different geographical locations can take advantage of SnapMirror and make local copies of mirrored data available to all locations, enhancing efficiency and productivity.

NetApp SnapVault

NetApp SnapVault leverages disk-based backup and block-level incrementals for reliable, low-overhead backup and recovery of NetApp storage, and is suitable for any environment.

With SnapVault, data protection occurs at the block level—copying only the data blocks that have changed since the last backup, not entire files. This enables backups to run more frequently and to use less capacity, because no redundant data is moved or stored.

For distributed organizations, this not only makes disk-based backup cost-effective, it offers the option of backing up directly from remote facilities to a core data center, centralizing management, and minimizing investment needs at the edge.

Monitoring and Management

Storage monitoring and management are critical to the success of the server environment. NetApp offers tools to monitor the health of storage systems, provide alerts, generate reports, and manage storage growth.

The OnCommand Management Software Portfolio

NetApp OnCommand management software is a family of products that improve storage and service efficiency by letting IT administrators control, automate, and analyze shared storage infrastructure.

OnCommand System Manager

System Manager is the simple yet powerful management tool for NetApp storage, and is easy to use for small to medium-sized businesses, and efficient for large enterprises and service providers.

With System Manager, you don't need to be a storage expert to manage NetApp storage systems. System Manager enables wizard-driven setup of aggregates, volumes, LUNS, qtrees, shares, and exports. It manages both NAS (CIFS, NFS) and SAN (iSCSI, FC) as a single tool, and provides a common look and feel. System Manager allows easy configuration of storage efficiency features like thin provisioning, data compression and deduplication.

System Manager supports integration with VMware ESX for virtual storage management, and despite its ease-of-use, it can be used to graphically manage advanced storage features like SnapMirror, SyncMirror, SnapLock, vFiler and Vservers. System Manager is also the means to manage the newest NetApp innovation—Data ONTAP 8.x operating in Cluster-Mode.

NetApp OnCommand Unified Manager

NetApp OnCommand Unified Manager monitors, manages, and generates reports on all of the NetApp storage systems in an organization. When you are using NetApp thin provisioning, NetApp recommends deploying Unified Manager and setting up email and pager notifications to the appropriate administrators. With thin-provisioned storage, it is very important to monitor the free space available in the aggregates. Proper notification of the available free space means that additional storage can be made available before the aggregate becomes completely full.



Reader Tip

For information about setting up notifications in Unified Manager, see the Configuring Alarms and Managing Aggregate Capacity section in the *Unified Manager Administration Guide* on NOW.

Unified Manager's policy-based management, global monitoring, and reporting allow you to automate your data protection operations.

Managing data protection can be complicated and time consuming. Most tools fail to give you a comprehensive and easy-to-understand view of your data protection environment. Additionally, they make it difficult to efficiently provision and use storage resources.

Unified Manager can simplify common data protection tasks and automate management across Snapshot, SnapMirror, SnapManager, SnapVault, and Open Systems SnapVault operations. It automates storage provisioning and provides global policy-based management, monitoring, and alerting.

Unified Manager makes it easy to define, apply, and update data protection policies across the enterprise. It minimizes effort, cuts administrative overhead, and helps to meet best practices and service-level agreements globally.

A simple dashboard shows comprehensive data-protection information at a glance, including unprotected data, alerts, and utilization.

Unified Manager automation combines with thin provisioning, deduplication, NetApp Snapshot, and block incremental technology to shrink the storage footprint and increase management efficiency.

NetApp Unified Manager can speed the creation of new NetApp storage resources and help improve capacity management of existing storage resources. Storage administrators can use Unified Manager's policy-based automation to create repeatable, automated provisioning processes to improve the availability of data, and enable provisioned storage to comply with policies. These processes are faster than manually provisioning storage, are easier to maintain than scripts, and help to minimize the risk of data loss due to misconfigured storage.

Unified Manager applies user-defined policies to consistently select the appropriate resources for each provisioning activity. This frees administrators from the headache of searching for available space to provision and allows more time for strategic issues. A centralized management console allows administrators to monitor the status of their provisioned storage resources.

Unified Manager can help improve your business agility and capacity utilization, shrink provisioning time, and improve administrator productivity. By leveraging Unified Manager's thin provisioning and deduplication capabilities, you can get a high level of storage efficiency from your NetApp storage investment. This allows you to store more data more efficiently, and helps improve your business agility.

NetApp SnapManager

NetApp SnapManager management tools integrate with the leading business applications to automate and simplify the complex, manual, and time-consuming processes associated with the backup, restoration, recovery, and cloning of the leading business applications, including Oracle, Microsoft Exchange, SQL Server, SharePoint, SAP, and server virtualization.

With NetApp SnapManager, you can:

- Leverage the NetApp technology stack to create near-instant and space-efficient Snapshot copies and clones of your applications.
- Integrate with native application technologies and achieve complete automation of data management.
- Use policies to simplify, standardize, and automate data protection.
- Increase back-up frequency—without affecting performance—for better data protection.
- Recover and restore a failed database to full production in minutes, regardless of size.
- Create complete data clones in seconds on primary storage or directly to your development and test environment.
- Use clones to engage in parallel QA, development, testing, and other processes, and deploy applications faster than ever before.

For more information on NetApp products, services, and solutions, NetApp sales representatives and reseller partners are ready to answer your questions and provide you with the pricing and configuration information you need to make your purchasing decision.

In the United States you can reach NetApp directly at: 1-877-263-8277 or select from the list of reseller partners located at:
<http://www.netapp.com/us/how-to-buy/>

If you are calling from outside the United States, please select your country from the list located on the right side of the link provided above.

Appendix A: Changes

This appendix summarizes the changes to this guide since the previous Cisco SBA series.

- We updated the NetApp storage system to a NetApp FAS3240 running Data ONTAP 8.0.3 operating in 7-Mode. The initial system setup reflects the changes required for the newer platform. Most of the GUI-based configuration in this guide remains the same, allowing an easy migration from the older NetApp FAS3140 to the NetApp FAS3240 storage system.
- We have added the SNMP, Network Time Protocol (NTP), and Cisco Discovery Protocol (CDP) options to the initial system setup.
- We made minor changes to improve the readability of this guide.

Notes

Feedback

Please use the [feedback form](#) to send comments and suggestions about this guide.



SMART BUSINESS ARCHITECTURE



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2013 Cisco Systems, Inc. All rights reserved.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)