



**SANtricity® 11.4x Installing and Configuring for
Linux® Power Guide for Advanced Users**

StorNext QD7000

Firmware 8.40.xx.xx



SANtricity® 11.4x Installing and Configuring for Linux® Power Guide for Advanced Users, 6-68662-01 Rev A, March 2018
Product of USA.

Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2018 Quantum Corporation. All rights reserved.

Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.

TRAEMARK STATEMENT

Artico, Be Certain (and the Q brackets design), DLT, DXi, DXi Accent, DXi V1000, DXi V2000, DXi V4000, GoVault, Lattus, NDX, the Q logo, the Q Quantum logo, Q-Cloud, Quantum (and the Q brackets design), the Quantum logo, Quantum Be Certain (and the Q brackets design), Quantum Vision, Scalar, StorageCare, StorNext, SuperLoader, Symform, the Symform logo (and design), vmPRO, and Xcellis are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

Products mentioned herein are for identification purposes only and may be registered trademarks or trademarks of their respective companies. All other brand names or trademarks are the property of their respective owners.

Quantum specifications are subject to change.



Preface

Note: The 8.40.xx.xx firmware (Madrid) is used in the QD7000 (E5600, Titan RAID controller, only). Refer to the [NetApp to Quantum Naming Decoder](#) section for additional information.

This section provides the following information:

- [Audience](#)
- [Prerequisites](#)
- [NetApp to Quantum Naming Decoder](#)
- [Product Safety Statements](#)
- [Contacts](#)
- [Comments](#)
- [Quantum Global Services](#)

Audience

This manual is intended for storage customers and technicians.

Prerequisites

Prerequisites for installing and using this product include knowledge of:




- Servers and computer networks
- Network administration
- Storage system installation and configuration
- Storage area network (SAN) management and direct attach storage (DAS)
- Fibre Channel (FC) and Ethernet protocols



NetApp to Quantum Naming Decoder

Use [Table 1](#) to correlate the NetApp product nomenclature to the equivalent Quantum-storage naming conventions.

Table 1 Product Nomenclature

E-Series NetApp Product	Quantum-Storage	Description
Controller-Drive Tray	Base System	Quantum uses Base System when referring to a drive tray with the RAID controllers.
Drive Tray	Expansion Unit	Quantum uses Expansion Unit when referring to a drive tray with the environmental services modules (ESMs).
E5600 (Code Name: Titan)	RAID Controller	Four 16Gb/s FC SFP+ host ports
E5500 (Code Name: Soyuz)	RAID Controller	Four 16Gb/s FC SFP+ host ports
E5400 (Code Name: Pikes Peak)	RAID Controller	Four 8Gb/s FC SFP+ host ports
DE6600 (Code Name: Wembley)	4U 60-drive enclosure	Sixty 3.5 inch disk drives

E-Series NetApp Product	Quantum-Storage	Description
<p>E5660</p> <ul style="list-style-type: none"> • DE6600 4U drive enclosure • With E5600 RAID controllers (Titan) 	<p>Quantum StorNext QD7000</p>	
<p>E5560</p> <ul style="list-style-type: none"> • DE6600 4U drive enclosure • With E5500 RAID controllers (Soyuz) 	<p>Quantum StorNext QD7000</p>	
<p>E5460</p> <ul style="list-style-type: none"> • DE6600 4U drive enclosure • With E5400 RAID controllers (Pikes Peak) 	<p>Quantum StorNext QD6000</p>	

E-Series NetApp Product	Quantum-Storage	Description
<p>E5424</p> <ul style="list-style-type: none"> • DE5600 24-drive 2U drive enclosure • Code Name: Camden • With E5400 RAID controllers (Pikes Peak) 	<p>Quantum StorNext QS2400</p>	
<p>E5412</p> <ul style="list-style-type: none"> • DE1600 12-drive 2U drive enclosure • Code Name: Ebbets • With E5400 RAID controllers (Pikes Peak) 	<p>Quantum StorNext QS1200</p>	

Product Safety Statements

Quantum will not be held liable for damage arising from unauthorized use of the product. The user assumes all risk in this aspect.

This unit is engineered and manufactured to meet all safety and regulatory requirements. Be aware that improper use may result in bodily injury, damage to the equipment, or interference with other equipment.

WARNING: Before operating this product, read all instructions and warnings in this document and in the system, safety, and regulatory guide.

警告 在使用本产品之前，请先阅读本文件及系统、安全和法规信息指南中所有的说明和警告信息。

警告 操作本产品前，请先阅读本文件及系统、安全与法规资讯指南中的指示与警告说明。

ADVERSAL Læs alle instruktioner og advarsler i dette dokument og i *Vejledning om system-sikkerheds- og lovgivningsoplysninger*, før produktet betjenes.

AVERTISSEMENT Avant d'utiliser ce produit, lisez la totalité des instructions et avertissements de ce document et du *Guide d'informations sur le système, la sécurité et la réglementation*.

HINWEIS Lesen Sie vor der Verwendung dieses Produkts alle Anweisungen und Warnhinweise in diesem Dokument und im System, Safety, and Regulatory Information Guide (Info-Handbuch: System, Sicherheit und Richtlinien).

אזהרה לפני ההפעלה של מוצר זה, קרא את כל ההוראות והאזהרות הכלולות במסמך זה וכן במדריך מידע בנושאי מערכת, בטיחות ותקינה

警告 この製品を使用する前に、本文書、および『システム、安全、規制に関する情報ガイド』に記載しているすべての警告と指示をお読みください。

경고 이 제품을 작동하기 전에 이 문서 및 시스템, 안전, 및 규제 정보 안내서에 수록된 모든 지침과 경고 표지를 숙지하십시오.

ПРЕДУПРЕЖДЕНИЕ

Перед началом эксплуатации данного устройства ознакомьтесь во всеми инструкциями и предупреждениями, приведенными в данном документе и в *Справочном руководстве по устройству, технике безопасности и действующим нормативам*.

ADVERTENCIA

Antes de utilizar este producto, lea todas las instrucciones y advertencias en este documento y en la Guía informativa sobre sistema, seguridad y normas.

WARNING

Läs alla anvisningar och varningar i detta dokument och i *System, säkerhet och krav från myndigheter - Informationshandbok* innan denna produkt tas i bruk.

Contacts

For information about contacting Quantum, including Quantum office locations, go to:

<http://www.quantum.com/aboutus/contactus/index.aspx>

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com

Quantum Global Services

Accelerate service issue resolution with these exclusive Quantum StorageCare services:

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Benefit today at:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>

- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can

also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Get started at:

<http://www.quantum.com/customercenter/>

For further assistance, or if training is desired, contact the Quantum Customer Support Center:

North America	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49-6131-324-185
APAC	+800-7826-8887 (toll free) +603-7953-3010

For worldwide support:

<http://www.quantum.com/serviceandsupport/get-help/index.aspx#contact-support>



SANtricity® 11.40

Installing and Configuring for Linux®

Power Guide for Advanced Users

September 2017 | 215-11889_A0
doccomments@netapp.com

 **NetApp®**

Contents

Deciding whether to use this Power Guide	5
Configuration options	6
Configuration worksheet	10
Deciding on the management method	12
Management methods	12
Out-of-band and in-band requirements	13
Installing SANtricity Storage Manager	16
Installing the storage array as a boot device	16
Installing SANtricity Storage Manager packages using silent mode	17
Deciding which packages to install	17
Host operating systems	17
Storage management software components	18
Installing the SANtricity software on hosts, monitors, and management stations	19
Adding the storage array to the management domain	25
Preparing to add the storage array to the management domain	25
Completing preliminary tasks for preparing the storage array	25
Setting IP addresses	25
Naming the storage array	26
Passwords	27
Choosing the method for adding the storage array to the management domain	28
Configuring management port IP addresses using the Quick Connect utility	29
Using automatic discovery to add storage arrays to the management domain	29
Manually configuring the controllers by setting up a temporary private network	30
Configuring management port using System Manager	31
Configuring a management port using Storage Manager	33
Configuring multipath	35
Overview of multipath drivers	35
Multipath driver setup considerations	35
Supported multipath drivers	36
Automatic Load Balancing feature overview	37
Multipath configuration diagrams	38
How a multipath driver responds to a data path failure	41
User responses to a data path failure	41
Failover drivers for the Linux operating system	42
Power methods for configuring multipath	42
Dividing I/O activity between two RAID controllers to obtain the best performance	42

Installing DM-MP	43
Overview of migrating to the Linux DM-MP multipath driver	44
Verifying correct operational mode for ALUA	48
Verifying OS compatibility with the Automatic Load Balancing feature	49
Migrating an existing DM-MP system to use Automatic Load Balancing ...	50
Setting up the multipath.conf file	50
Setting up DM-MP for large I/O blocks	53
Using the device mapper devices	54
Rescanning devices with the DM-MP multipath driver	55
Troubleshooting Device Mapper	56
Configuring host utilities, virtualization, and clustering	57
Host utilities	57
Virtualization considerations	58
Multipathing and virtualization	59
Linux-specific virtualization options	59
Host clustering support	59
Cluster accessibility	60
Cluster topology	61
Cluster shared storage in SANtricity	61
What are SCSI reservations?	62
Deciding whether to use disk pools or volume groups	63
Creating a volume group	63
Creating a volume group using the AMW	66
Host-side storage considerations	67
Storage partitions	67
Using the iscsiadm open-iscsi administration utility with E-Series products	70
Linux partitions, filesystems, and mounting recommendations	72
Services on Linux that apply to E-Series products	73
Copyright information	74
Trademark information	75
How to send comments about documentation and receive update notifications	76

Deciding whether to use this Power Guide

You can customize the installation and configuration of the management software and E-Series storage array to fit your data storage requirements. The quickest path is described in the SANtricity Express Guide for your operating system. This Power Guide provides additional options beyond those included in the Express Guides. You can use a mixture of express methods and power methods to customize your installation.

Use this document for one of the following reasons:

You have...	...and you want to...
Planned for an express installation of SANtricity Storage Manager or an express configuration of SANtricity System Manager on your operating system	<ol style="list-style-type: none"> 1. Review the options for managing your storage array by exploring the table of contents of the Express Guide and this Power Guide. 2. Verify your decisions by using the Configuration worksheet on page 10. 3. Proceed through the Express Guide for your operating system. Review the options in this Power Guide and choose the variations you want to consider for your storage installation.
Completed an express method install using one of the E-Series Express Guides	Review the options for managing your storage arrays. See Configuration options on page 6.
An active E-Series configuration	Consider adding options or modifying your installation: <ol style="list-style-type: none"> 1. Verify your decisions by using the Configuration worksheet on page 10. 2. Read the conceptual information and optional procedures in this Power Guide. 3. Follow the procedures that are appropriate for your data storage requirements.

Related information

[NetApp E-Series Systems Documentation Center](#)

Configuration options

When planning the installation of an E-Series storage array, you can consider a number of options beyond the express method, including how to install the storage management software, how to manage the domain, and how to configure AutoSupport and alerts.

Type of storage array

If you have E-Series storage arrays, you could have one or more of these models:

- E5700
- E2800
- E5600
- E2700

Your options for storage management software vary depending on the array type.

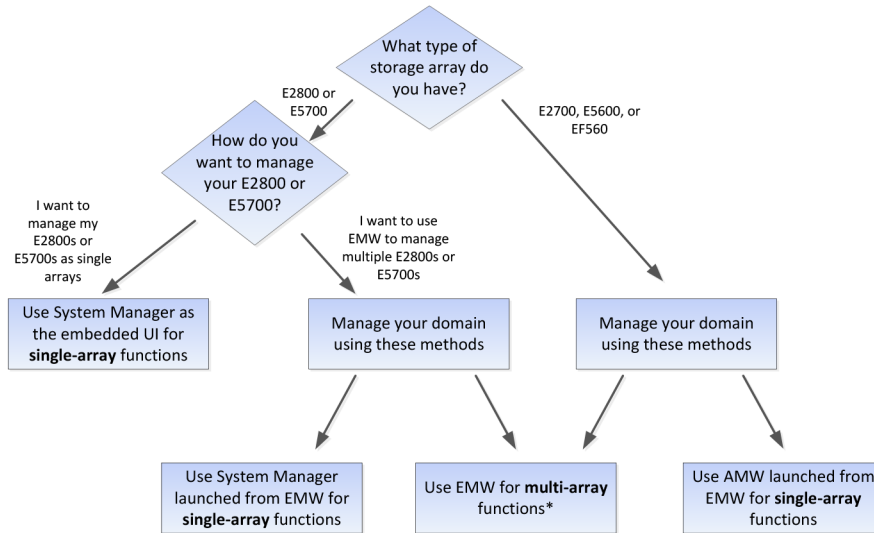
Storage management software

NetApp's two software interfaces, SANtricity **Storage** Manager and SANtricity **System** Manager, are each appropriate in specific use cases:

- SANtricity Storage Manager is compatible with the E2700 and E5600 storage arrays. SANtricity Storage Manager's client-based user interface has an **Enterprise Management Window (EMW)** and an **Array Management Window (AMW)**.
 - The EMW provides functions for configuring and managing multiple arrays.
 - The AMW provides functions for configuring and managing a single array. You launch the AMW from within the EMW.
- SANtricity System Manager's browser-based user interface is appropriate for managing either single or multiple E2800 or E5700 arrays. How you launch SANtricity Storage Manager depends on whether you want to manage a single array or multiple arrays.
 - To manage one or more E2800 or E5700 arrays as single arrays, launch System Manager in a browser.
 - To manage one or more E2800 or E5700 arrays as a multiple-array configuration, launch System Manager from the EMW.


Important: To launch System Manager from the EMW, you must use a desktop environment (examples: KDE, GNOME). If you are not using a desktop environment, you can open your browser manually and point to the array to open System Manager.

Use the following decision tree to help you determine which storage management software you will use.

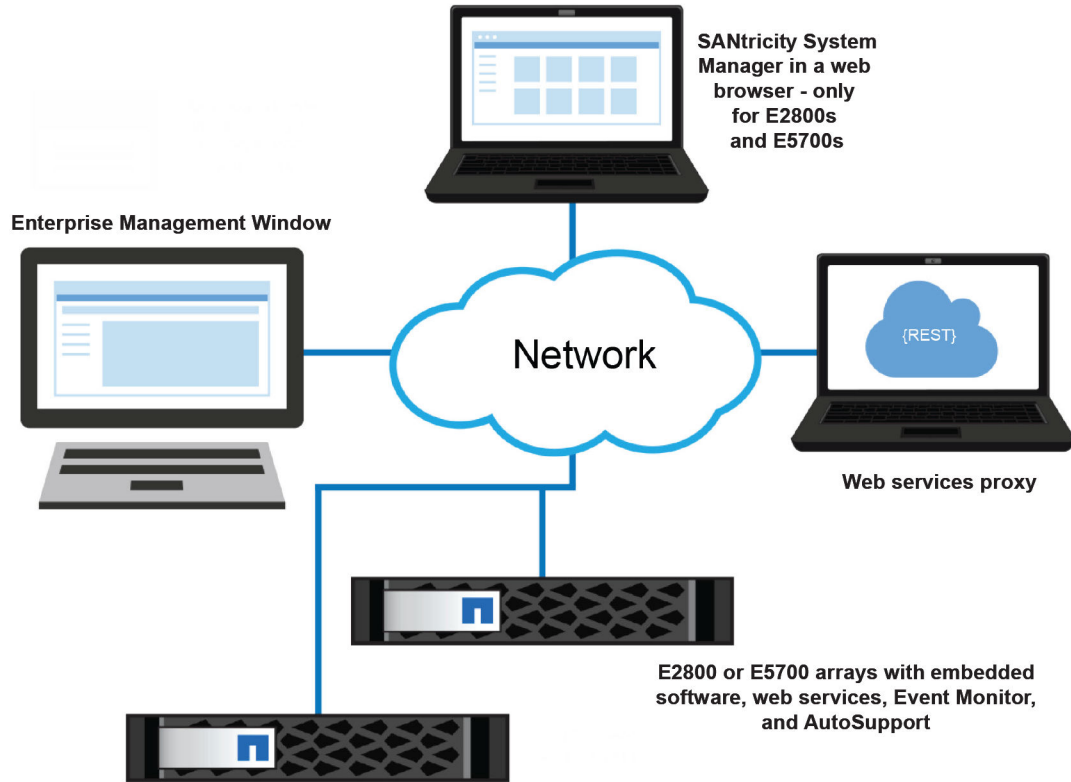


The following configuration examples further illustrate the use of the appropriate storage management software.

- Single E2800 or E5700 storage array** — If you have a single E2800 or E5700 array and are not using either the Synchronous Mirroring or Asynchronous Mirroring feature, all configuration can be handled from SANtricity System Manager. You can perform a host install of Storage Manager to get the host context agent (SMagent) to pre-populate host information in SANtricity System Manager. For more information about host installations, refer to [Installing the SANtricity software on hosts, monitors, and management stations](#) on page 19.




- Multiple E2800 or E5700 storage arrays** — If you have more than one E2800 or E5700 storage array, you can install the EMW to manage your storage environment while handling storage array-based configuration through SANtricity System Manager. The EMW is included with SANtricity Storage Manager.



Note: If you are not using Synchronous or Asynchronous Mirroring features, you do not need to install the EMW. Instead, you can bookmark multiple SANtricity System Manager storage arrays in a browser.

- **Mixed array environment** — You must use the EMW that is part of the SANtricity Storage Manager installation if either of the following statements is true:
 - You have one or more E2800 storage arrays and any E2700, E5600, or EF560 storage arrays and want to have the E2800 or E5700 storage array included in your aggregate view.
 - You want to use Synchronous or Asynchronous Mirroring.

For array-based tasks on the E2800 or E5700 storage arrays, use SANtricity System Manager launched from the EMW. For array-based tasks on E2700, E5600, or EF560 storage arrays, use the AMW launched from the EMW.



AutoSupport and alerts

You configure AutoSupport (ASUP), email, and syslog alerts differently, depending on the type of storage array:

- **E2800, E5700** — You must configure AutoSupport and alerts on each storage array. These components are embedded in the E2800 and E5700 controllers.
- **E2700, E5600, and EF560** — You can configure AutoSupport and alerts globally by using the EMW.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Configuration worksheet

The storage configuration worksheet allows you to track your decisions about your E-Series configuration. Express methods and power methods are listed.

Circle your components and options in the table. For express method instructions, see the Express Guide for your operating system (OS).

Decision/Component	Express method	Power method (described in this Power Guide)
Controller model	<ul style="list-style-type: none"> • E5700 • E2800 • E2700 • E5600 • EF560 	<ul style="list-style-type: none"> • E5700 • E2800 • E2700 • E5600 • EF560 <p>See Configuration options on page 6.</p>
Storage management method (physical connectivity)	Out-of-band	In-band See Deciding on the management method on page 12.
<p>Management software components</p> <p>You use SANtricity Storage Manager or SANtricity System Manager for different storage arrays and different purposes. See Configuration options on page 6.</p>	<ul style="list-style-type: none"> • SANtricity Storage Manager <ul style="list-style-type: none"> ◦ Enterprise Management Window (EMW) ◦ Array Management Window (AMW) ◦ CLI ◦ Event Monitor • SANtricity System Manager <ul style="list-style-type: none"> ◦ For E2800 or E5700 controller shelves ◦ Not a separate installation ◦ Browser-based • Multipath driver • Unified Host Utilities 	<ul style="list-style-type: none"> • SMagent (part of the host manager installation) • Multipath driver • Other utilities, such as SMdevices <p>See Deciding which packages to install on page 17.</p>
Using the storage array as a boot device	No	Yes See Installing the storage array as a boot device on page 16.

Decision/Component	Express method	Power method (described in this Power Guide)
Using Silent Mode when installing SANtricity Storage Manager	No	Yes See <i>Installing SANtricity Storage Manager packages using silent mode</i> on page 17.
I/O protocol	All protocol-specific tasks are described in Express Guides	No additional protocol-specific options.
Management IP addressing method	Static IP, using Quick Connect utility	<ul style="list-style-type: none"> • Static IP, by temporarily setting up a private network • Dynamic host configuration protocol (DHCP) • IPv6 stateless address auto configuration See <i>Setting IP addresses</i> on page 25 and <i>Choosing the method for adding the storage array to the management domain</i> on page 28.
Disk pools (pools) or volume groups	Disk pools (pools)	Disk pools (pools) or volume groups See <i>Deciding whether to use disk pools or volume groups</i> on page 63.

Related references

Configuration options on page 6

Deciding on the management method

Before you install and use either SANtricity System Manager software or SANtricity Storage Manager software, you need to know which storage management method you plan to use.

Management methods

You can choose the best management method based on your system configuration and management goals. You manage a storage array from a management station or from a host.

Management methods include:

- Out-of-band management
- In-band management
- A combination of out-of-band and in-band management

Storage management includes these activities:

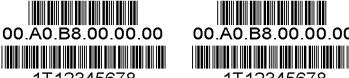
- Configuring available storage array capacity to maximize data availability, optimize application performance, and make the most of storage resources
- Configuring destinations to receive alert messages for critical problems concerning one or more storage arrays
- Monitoring storage arrays for problems or conditions that require attention
- Recovering from storage array problems to maximize data availability

Note: For E2800 and E5700 controllers, the in-band management method is supported only through the CLI.

Out-of-band and in-band requirements

To determine whether to use out-of-band or in-band management, consider the requirements, advantages, and disadvantages of each method.

Management method	Requirements	Advantages	Disadvantages
All out-of-band methods	Connect separate Ethernet cables to each controller.	<p>This method does not use a logical unit number (LUN) on the host.</p> <p>This method does not use I/O path bandwidth for storage array management functions.</p> <p>You do not need to install host-agent (SMagent) software.</p> <p>This method does not use the SAS, Fibre Channel or iSCSI bandwidth for storage array management functions.</p>	<p>Ethernet cables are required.</p> <p>Does not allow you to choose which controller is used for the EMW. Controller A is used until SANtricity Storage Manager has difficulty communicating on that path. Then the system switches to controller B.</p>
Out-of-band <i>without</i> a DHCP server	Manually configure the network settings on the controllers.	--	You must manually configure the network settings on the controllers.
Out-of-band – IPv6 stateless address auto-configuration <i>without</i> a DHCP server (IPv6 networks only)	<p>Connect at least one router for sending the IPv6 network address prefix in the form of router advertisements.</p> <p>The router is necessary to route the IPv6 packets outside the local network.</p>	<p>No additional manual network configuration is required on the controllers.</p> <p>By default, the controllers automatically obtain their IP addresses by combining the auto-generated link local address and the IPv6 network address prefix after you turn on the power to the controller-drive tray.</p>	A router is required.

Management method	Requirements	Advantages	Disadvantages
<p>Out-of-band <i>with</i> a DHCP server (IPv4 networks only)</p>	<p>Connect separate Ethernet cables to each controller.</p> <p>Assign either static IP addresses or dynamic IP addresses to the controllers using your DHCP server. Alternatively, both the SANtricity System Manager and the SANtricity Storage Manager AMW can be used to set the IP addresses after the storage array has been discovered. It is recommended that you either reserve the controller IPs in the DHCP server or assign a static IP address so that the management port addresses will not change if the power to the storage array is disrupted.</p> <p>Check your DHCP server for the IP addresses that are associated with the media access control (MAC) addresses of the controllers.</p> <p>The MAC address appears on a label on each controller in the form: <i>xx.xx.xx.xx.xx.xx</i> .</p> <div data-bbox="483 1192 863 1289" style="border: 1px solid black; padding: 5px; text-align: center;">  <p>00.A0.B8.00.00.00 00.A0.B8.00.00.00</p> <p>1T12345678 1T12345678</p> </div>	<p>No additional manual network configuration is required on the controllers.</p> <p>By default, the controllers automatically obtain their IP addresses from the DHCP server after you turn on the power to the controller-drive tray.</p> <p>This method does not use a special Access Volume to communicate with the host.</p>	<p>No additional disadvantages.</p>

Management method	Requirements	Advantages	Disadvantages
In-band	<p>Install the host agent software (SMagent) on at least one of the I/O-attached hosts. (To locate the SMagent, refer to Storage management software components on page 18.)</p> <p>The host-agent software, which is included with the storage management software, manages the storage array through the data path from an I/O-attached host or an Ethernet connection from a storage management station to the I/O-attached host that is running the host-agent software.</p> <p>The in-band method requires a special access volume to communicate between the host and the storage array. This volume is created automatically.</p> <p>If a firewall is installed on the I/O-attached host, ensure that port 2463 is open.</p>	No additional manual network configuration is required on the controller.	<p>This method:</p> <ul style="list-style-type: none"> • Uses both a LUN on the host and the SAS, Fibre Channel, or iSCSI bandwidth for storage array management functions. • Is not supported on System Manager; you must use the CLI. • Does not allow you to choose which controller is used for the command-line interface (SMcli).

Installing SANtricity Storage Manager

If the express method of installing SANtricity Storage Manager does not meet the requirements of your configuration, you can consider alternate power methods. These methods apply to Storage Manager only, and not System Manager. System Manager is embedded in the controller, so you do not need to install it.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Installing the storage array as a boot device

Before you install the storage management software components on the host, you must prepare the storage array and the host. Because E-Series storage behaves as a block device, you can install an operating system on it and boot that operating system from an E-Series storage array, instead of relying on local storage.

Using the E-Series storage array as a boot device serves as a less expensive, potentially faster alternative to internal storage. For example, if operating a Blade system, this process is much less expensive than purchasing internal storage for all blades. This process is called SAN booting - or relying on the SAN to boot a host. The concept of SAN boot is straight forward; however, the execution can become complicated.

The following describes the overall workflow required for setting up a SAN boot on E-Series storage:

- The host, and more specifically the adapter attached to E-Series storage, is directed to present a mapped or assigned volume from storage prior to boot (in BIOS, uEFI, or another appropriate type of firmware).
This process is vendor-specific, protocol-specific, and architecture specific.
- The host can boot using the installation media.
- The installation selects the volume provided by storage to install.
Sometimes this requires a driver update disk (DUD). Additionally, failover might or might not have to be loaded during this step, depending on the operating system.
- After reboot, the boot options must set the newly-installed volume as the primary boot option.
This step is vendor-specific for the adapter as well as the server.

Note: NetApp recommends using LUN 0 for booting, and some operating systems might require it.

Boot device support

Not all operating systems support the use of a storage array as a boot device. Support for using a boot device also depends on the type of host connection. For example, Fibre Channel and SAS connections are supported, while iSER over Infiniband, SRP over InfiniBand and some iSCSI connections are not supported.

The following table shows which operating systems support this configuration, but you should consult the *[Interoperability Matrix Tool](#)* to ensure that your HBA and operating system are supported.

Operating system	Boot device support	Recommended number of paths for installation
AIX	Yes, where supported by the HBAs	2
HP-UX	Yes, where supported by the HBAs	2
Linux	Yes, where supported by the HBAs	2
Mac OS X	No	1
Solaris	Yes, where supported by the HBAs	2
VMware	Yes, where supported by the HBAs	2
Windows	Yes, where supported by the HBAs	1 (works with 2, but 1 is recommended)

Installing SANtricity Storage Manager packages using silent mode

You can use the Silent installation mode for any OS that is supported by Install. Silent mode requires minimal user interactions and is useful when deploying a large number of servers that are not connected to terminals.

To install the storage manager packages using the Silent mode, locate the specified components in the `installer.properties` file by entering the following command for your operating system:

Linux/Unix: `SMIA.xx.xx.xx.xx.bin -f installer.properties`

This command creates the `installer.properties`.

Deciding which packages to install

Different storage management software components and packages are required for different machines. Additionally, you will install different components depending on the environment you need to support for your particular configuration.

Host operating systems

Considerations for both SANtricity System Manager and SANtricity Storage Manager's support of host operating systems (OSes) include OS versions, host bus adapters (HBAs), host processors, multipath drivers, JRE levels, and SANboot.

For information about compatibility of these components with SANtricity Storage Manager, see the [NetApp Interoperability Matrix Tool](#).

Storage management software components

Depending on your configuration and data storage requirements, you select different storage management software components.

SANtricity Storage Manager or SANtricity System Manager?

To configure and manage E2700 or E5600 storage arrays, you use SANtricity Storage Manager's Array Management Window (AMW) and Enterprise Management Window (EMW). If you have an E2800 or E5700 storage array, you configure it using the browser-based SANtricity System Manager rather than through SANtricity Storage Manager's AMW. If you have multiple types of storage arrays or more than one E2800 or E5700 and want to manage your entire environment, you install and use SANtricity Storage Manager's EMW.

SANtricity System Manager is browser-based, so there is no installation required. After you install your E2800 or E5700 hardware and connect it to the network by assigning appropriate IPs, subnet masks, and the gateway for the controllers, you access SANtricity System Manager by pointing a browser to the E2800 or E5700's IP address or domain name.

SANtricity Storage Manager components

Client

This package contains both the graphical user interface (GUI) (containing both the EMW and the AMW) and the command line interface (CLI) for managing the storage arrays. This package also contains the Event Monitor that sends alerts when a critical problem exists with the storage array.

Utilities

This package contains utilities that let the operating system recognize the volumes that you create on the storage array and to view the operating system-specific device names for each volume.

Agent

This component contains software that allows a management station to communicate with the controllers in the storage array over the I/O path of a host (see *Out-of-band and in-band requirements* on page 13). This package is required for in-band management, and can be used for out-of-band as well to pre-populate host port information on all data hosts for both AMW and SANtricity System Manager.

Multipath driver

For this operating system, the preferred multipath driver is included "in-box" and is not available from the SANtricity host install package. The multipath driver manages the I/O paths into the controllers in the storage array. If a problem exists on the path or a failure occurs on one of the controllers, the driver automatically reroutes the request from the hosts to the other controller in the storage array. Always check the *Interoperability Matrix Tool* to verify what multipath drivers are supported for your configuration.

You must install the utilities and the multipath driver on each host attached to the storage array.

Hosts

The host adapters in the hosts that are attached to the storage array are known to the storage management software. However, in most cases the storage management software does not know which host adapters are associated with which hosts. Only when the SMagent services runs on the host that is attached to a storage array can the storage management software associate HBA ports to that host.

Note: If your operating system configures automatically, then, by default, the host context agent automatically defines all attached hosts that are running SMagent in the

mapping view of the AMW with a default mapping scheme which you can modify to the needs of your configuration.

Event Monitor

During the client installation, you might be asked whether you want to start the Event Monitor.

If you are running an E2800 or E5700 storage array, the Event Monitor resides on the controller and must be configured for each storage array. Use either SANtricity System Manager or the SMcli to complete the configuration task.

If you have an E2700 or E5600 storage array, start the monitor on only one management station that runs continuously. If you start the monitor on more than one management station, you receive duplicate alert notifications about problems with the storage array. If you install SANtricity components on more than one management station and are not asked about the Event Monitor, verify that the monitor is active on only one of the systems.

Note: To receive critical alert notifications and to access the AutoSupport (ASUP) feature with E2700 or E5600 storage arrays, you must have Event Monitor running on just one management station. With the E2800 or E5700 storage array, AutoSupport functionality is embedded in the controller.

Related information

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Installing the SANtricity software on hosts, monitors, and management stations

You can use the following software configuration diagrams and accompanying tables to determine which software packages to install on each machine (host, monitor, or management station):

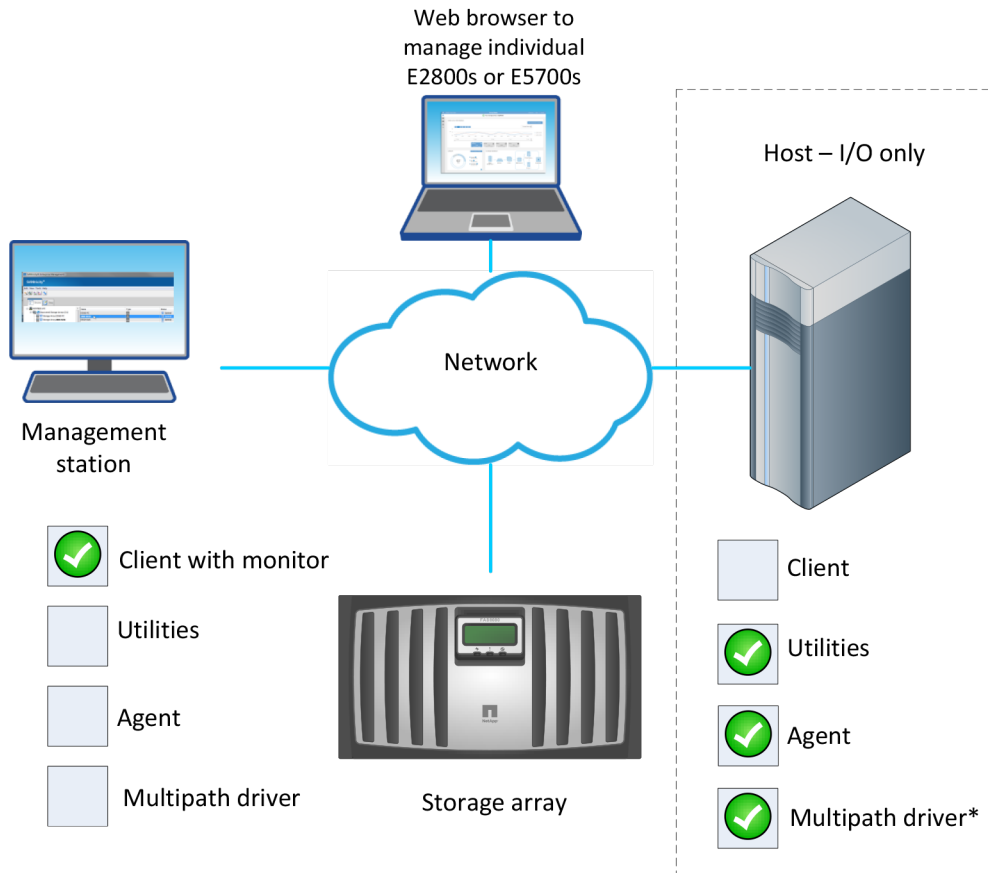
For this operating system, the preferred multipath driver is included "in-box" and is not available from the SANtricity host install package

The following table shows the packages that apply to particular installations.

Installation wizard selections			
Type of installation	Client	Utilities	Agent
Typical Installation	✓	✓	✓
Management Station	✓	—	—
Host	—	✓	✓
Custom (you select the components)	✓	✓	✓

Installing on the host (I/O only)


The following conceptual diagram and table provide basic information for installing on the host only for I/O.



Machines and required software: Host (I/O only)		
Minimum Software Required	Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
<ul style="list-style-type: none"> • Utilities • Agent • Multipath driver* 	Host	
*The preferred Multipath driver is "in-box" with the operating system for Linux, Solaris, and VMware systems.		

Installing Host -- Also acting as an agent for the in-band management method


The following conceptual diagram and table provide basic information for installing the host for in-band management.



Machines and required software: Host -- Also acting as an agent for the in-band management method	
Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
<ul style="list-style-type: none"> • Typical Installation • Host (no client install) • Custom 	Click No to the prompt, Automatically start Monitor?

Installing host also acting as monitor, and management stations

The following conceptual diagram and first table provides basic information for installing the host to act as a monitor for sending critical alerts. The management station installation options on a separate system are also included in the table that follows.



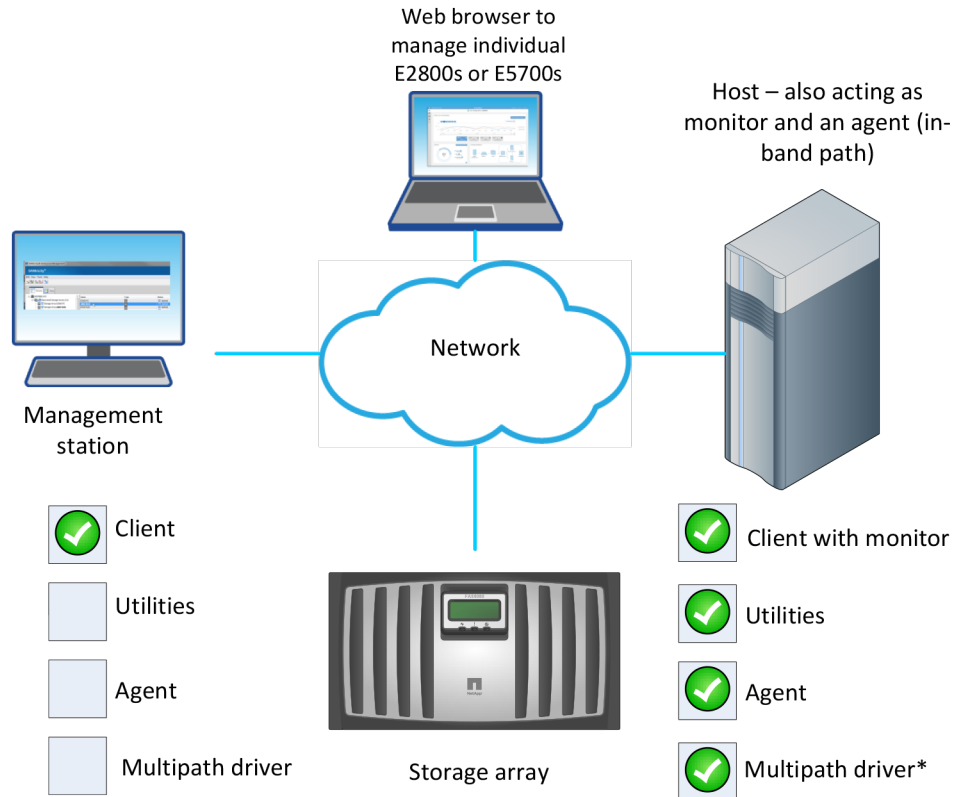
Machines and required software: Host as Monitor for sending critical alerts	
Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
<ul style="list-style-type: none"> • Typical Installation • Custom 	<ul style="list-style-type: none"> • Click Yes to the prompt, Automatically start Monitor? • Start the monitor on only one host that will run continuously.

Machines and required software: Management Station options			
Machine	Minimum Software Required	Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
Management station*	Client**	<ul style="list-style-type: none"> • Typical Installation • Management Station • Custom 	<ul style="list-style-type: none"> • Click No to the prompt, Automatically start Monitor?

Machines and required software: Management Station options			
Machine	Minimum Software Required	Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
Management station with the Storage Manager Event Monitor always running*	Client	<ul style="list-style-type: none"> • Typical Installation • Management Station • Custom 	<ul style="list-style-type: none"> • Click Yes to the prompt, Automatically start Monitor?
<p>*If you are managing a single E2800 or E5700 storage array, you do not need a separate Management station.</p> <p>**Linux OS automatically starts the Event Monitor when you install the storage management software, and it will continue to run until you stop it.</p>			

Installing host that acts as monitor and an agent (in-band management path)

The following conceptual diagram and table provide basic information for installing the host to act as a monitor for sending critical alerts and an agent for in-band management.



Machines and required software: Host also acting as monitor and an agent (in-band management path) and monitor for sending critical alerts		
Minimum Software Required	Installation Package (Choose One) (See the Installation wizard selections table above.)	Notes
<ul style="list-style-type: none"> • Client • Utilities • Agent • Multipath driver 	<ul style="list-style-type: none"> • Typical Installation • Custom 	<ul style="list-style-type: none"> • Click Yes to the prompt, Automatically start Monitor? • Start the monitor on only one host that will run continuously.

Adding the storage array to the management domain

Before you add the storage array to the management domain, review the guidelines and complete the preliminary tasks. Then, choose from a list of methods for adding the storage array.

Preparing to add the storage array to the management domain

You must prepare the storage array before adding it to the management domain, which consists of discovering any storage array within the local sub-network so that they display within the EMW.

Completing preliminary tasks for preparing the storage array

You complete some preliminary tasks before you can add the storage array to the management domain.

Make sure you have taken these steps:

- Connected all of the applicable cables.
- Turned on the power to the storage array (powering on the attached drive trays first, and then the controller-drive tray or controller tray).
- Installed the applicable storage management software.

Setting IP addresses

If the express method of using the Quick Connect utility to assign static IP addresses does not meet the requirements of your configuration, you can use one of the alternate methods for configuring IP addresses.

By default, E-Series controllers ship with DHCP enabled on both network ports. You can assign static IP addresses, use the default static IP addresses, or use DHCP-assigned IP addresses. You can also use IPv6 stateless auto-configuration.

Note: IPv6 is disabled by default on new E-Series systems, but you can configure the management port IP addresses using an alternate method, and then enable IPv6 on the management ports using SANtricity System Manager.

When the network port is in a "link down" state, that is, disconnected from a LAN, the SANtricity Storage Manager reports its configuration as either static, displaying an IP address of 0.0.0.0 (earlier releases), or DHCP enabled with no IP address reported (later releases). After the network port is in a "link up" state (that is, connected to a LAN), it attempts to obtain an IP address through DHCP.

If the controller is unable to obtain a DHCP address on a given network port, it reverts to a default IP address, which might take up to 3 minutes. The default IP addresses are as follows:

```
Controller 1 (port 1): IP Address: 192.168.128.101
```

```
Controller 1 (port 2): IP Address: 192.168.129.101
```

```
Controller 2 (port 1): IP Address: 192.168.128.102
```

```
Controller 2 (port 2): IP Address: 192.168.129.102
```

When assigning IP addresses:

- Reserve Port 2 on the controllers for Customer Support usage. Do not change the default network settings (DHCP enabled).
- To set static IP addresses for E2800 and E5700 controllers, use SANtricity System Manager. To set static IP addresses for E2700 and E5600 controllers, use SANtricity Storage Manager. After a static IP address is configured, it remains set through all link down/up events.
- To use DHCP to assign the IP address of the controller, connect the controller to a network that can process DHCP requests. Use a permanent DHCP lease.

Note: The default addresses are not persisted across link down events. When a network port on a controller is set to use DHCP, the controller attempts to obtain a DHCP address on every link up event, including cable insertions, reboots, and power cycles. Any time a DHCP attempt fails, the default static IP address for that port is used.

Related concepts

[Choosing the method for adding the storage array to the management domain](#) on page 28

Naming the storage array

You have some flexibility and some specific requirements when naming your storage array.

Take note of the following when naming your storage array:

- You can use letters, numbers, and the special characters underscore (_), hyphen (-), and pound sign (#). No other special characters are permitted.
- Limit the name to 30 characters. Any leading and trailing spaces in the name are deleted.
- Use a unique, meaningful name that is easy to understand and to remember. Avoid arbitrary names or names that would quickly lose their meaning in the future. The prefix “Storage Array” is automatically added to the name you assign. The full name is shown in the Logical pane and in the Enterprise Management Window. For example, if you named the storage array “Engineering,” it appears as “Storage Array Engineering.”
- The storage management software does not check for duplicate names. Check the Enterprise Management Window to make sure that the name you have chosen is not used by another storage array.
- When you first discover a storage array or manually add it, the storage array will have a default name of “unnamed.”

Passwords

Access Management, new in the 11.40 release, requires that users log in to SANtricity System Manager with assigned login credentials. Each user login is associated with a user profile that includes specific roles and access permissions. If you do not want to use Access Management, or if you have an E2700 or E5600 storage array for which the feature is unsupported, you can configure each storage array with an Administrator password. An optional Monitor password is available for E2700 and E5600 arrays.

Administrators can implement Access Management using one or both of these methods:

- Using RBAC (role-based access control) capabilities enforced in the storage array, which includes pre-defined users and roles.
- Connecting to an LDAP (Lightweight Directory Access Protocol) server and directory service, such as Microsoft's Active Directory, and then mapping the LDAP users to the storage array's embedded roles.

If you do not use Access Management or it is not supported, setting an Administrator password for your storage array protects it from being modified by unauthorized users. Modifying commands includes any functions that change the state of the storage array, such as creating volumes and modifying the cache settings. Setting a Monitor password allows users, who are not allowed to modify storage array configurations, to view storage array configurations and to monitor storage array health conditions.

Note that a Monitor password is not supported with SANtricity System Manager.

On SANtricity System Manager, you are asked if you want to set an Administrator password during initial set up.

On SANtricity Storage Manager, you are asked for a password only when you first attempt to change the configuration (such as creating a volume) or when you first perform a destructive operation (such as deleting a volume). You must exit both the Array Management Window and the Enterprise Management Window to be asked for the password again.

Follow these guidelines for setting passwords:

- For increased protection, use a long password with at least 15 alphanumeric characters. The maximum password length is 30 characters.
- Passwords are case sensitive.
- If you no longer want to have the storage array password-protected, enter the current password, and then leave the **New password** text box and the **Confirm password** text box blank.

Note: Only a user with the Administrator password can set or change the Monitor password. If a user with View-only access (Monitor Password) attempts to launch the Set Password dialog, the system prompts for the Administrator password.

Note: Both the Administrator storage array password and the Monitor storage array password are different from the pass phrase used for Drive Security.

Note: If you forget your password, you must contact your technical support representative for help to reset it.

Choosing the method for adding the storage array to the management domain

You can choose from several methods for adding the storage array to the management domain. The appropriate method depends on your network configuration and how you initially configured the controllers.

There are three primary methods of configuring the management ports of a storage array and adding them to the management domain:

- **Static IP addressing** - An Internet Protocol (IP) address for each management port that you enter. These addresses are typically assigned by a network administrator.
- **DHCP addressing** - An Internet Protocol (IP) address that the Dynamic Host Configuration Protocol (DHCP) server assigns. DHCP provides three mechanisms for IP address allocation. Automatic allocation is defined as DHCP assigning a permanent IP address to a client. Dynamic allocation is defined as DHCP assigning an IP address to a client for a limited time period or until the client explicitly lets go of the address. Manual allocation is defined as the network administrator assigning the IP address of the client, and DHCP conveys the assigned address to the client. A network uses one or more of these mechanisms, depending on the policies that the network administrator specifies.
- **Stateless address autoconfiguration with IPv6** - With stateless auto-configuration, hosts do not obtain addresses and other configuration information from a server. Stateless auto-configuration in IPv6 features link-local addresses, multicasting, and the Neighbor Discovery (ND) protocol. IPv6 can generate the interface ID of an address from the underlying data link layer address.

Note: You can change the configuration of a storage array to use a different type of management port IP addressing at any time. See the SANtricity System Manager online help or the SANtricity Storage Manager online help for detailed procedures.

Use one of the following methods to connect your E-Series storage arrays to the management domain:

If you are using...	...do this...
Static IP addressing	Use the Quick Connect utility. See Configuring IP addresses using the Quick Connect utility on page 29.
DHCP addressing of the management ports	Use auto-discovery to discover your storage array. See Using automatic discovery to add storage arrays to the management domain on page 29.
Stateless address auto-configuration, no DHCP server	Note: For E2700 and E5600 arrays only, the management station must reside on the same subnetwork as the array during controller management IP configuration.
Static IP addressing, and need an alternative to using the Quick Connect Utility	Temporarily set up a private network to configure the management ports. Note: For E2700 and E5600 arrays only, you will first need to configure the management station so that it resides on the same subnetwork during controller management IP configuration. See Manually configuring the controllers by setting up a temporary private network on page 30.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Configuring management port IP addresses using the Quick Connect utility

In this best-practices method for configuring communications, you configure the management station and array controllers to communicate using the Quick Connect utility.

Before you begin

- You have obtained the network configuration information from your network administrator for the controllers (IP address, subnet mask, and gateway or IP address and routable IP address).

About this task

The following figures show the location of management port 1 on the controllers.

Steps

1. Go to [SANtricity Quick Connect](#). Download and install the utility.
2. Follow the directions on the Wizard screens to configure your management port and to configure the IP address of each controller.
3. Connect an Ethernet cable to management port 1 (labeled P1) on each controller, and connect the other end to your network.

Note: Do not use port 2 on either controller. These ports are reserved for use by NetApp technical personnel.

Using automatic discovery to add storage arrays to the management domain

You can use automatic discovery to set the controller IP addresses using out-of-band management.

Before you begin

- The management station must be attached to the same subnet as the storage.
- Ethernet cables must be attached to each controller.
- The DHCP server must be configured to assign a permanent (static) DHCP lease.
- If you are using IPv6 stateless address auto configuration without a DHCP server, you must have connected at least one router for sending the IPv6 network address prefix in the form of router advertisements. By default, the controllers automatically obtain their IP addresses by combining the auto-generated link local address and the IPv6 network address prefix after you turn on the power to the controller-drive tray.

About this task

This procedure specifically applies to users with SANtricity Storage Manager configurations. If you have a SANtricity System Manager configuration, refer to [Configuring management port using System Manager](#) on page 31.

Steps

1. Open **SANtricity Storage Manager**.

The **Enterprise Management Window (EMW)** is displayed.

2. On the **Select Addition Method** screen, select the **Automatic** radio button, and then select **OK**.
This process finds all of the storage arrays on the local sub-network. Several minutes might lapse to complete the process.
3. Name the storage array.
 - a. In the **EMW Setup** tab, select **Name/Rename Storage Arrays**.
 - b. In the **Select storage array** list, select the storage array you added.
 - c. In the **Storage array name** field, type a name for the storage array.

Storage array names must not exceed 30 characters and cannot contain spaces. Names can contain letters, numbers, underscores (_), hyphens(-), and pound signs (#). Choose a descriptive name for the storage array to make it easier for data center administrators to manage the storage resources over time.

Manually configuring the controllers by setting up a temporary private network

You can manually configure the IP addresses on the controllers by setting up a temporary private network.

Before you begin

- You have connected the management station directly into Ethernet port 1 on each controller.
- You have connected an ethernet cable to the management station and to the management port 1 on A.

Note: Do not use port 2 on either controller. These ports are reserved for use by NetApp technical personnel.

- You have obtained the network configuration information from your network administrator for the controllers (IP address, subnet mask, and gateway or IP address and routable IP address).

Note: All controller shelves use Auto-MDIX (automatic medium-dependent interface crossover) technology to detect the cable type and configure the connection to the management station accordingly.

Steps

1. Change the IP address on the TCP/IP port on the management station from an automatic assignment to a manual assignment by using the default IP address subnet of the controllers.
 - a. Make note of the current IP address of the management station so that you can revert back to it after you have completed the procedure.

Note: You must set the IP address for the management station to something other than the controller IP addresses (for example, use 192.168.128.100 for an IPv4 network, or use FE80:0000:0000:0000:02A0:B8FF:FE29:1D7C for an IPv6 network).

Note: In an IPv4 network, the default IP addresses for Ethernet port 1 on controller A and controller B are 192.168.128.101 and 192.168.128.102, respectively.

- b. Change the IP address. Refer to your operating system documentation for instructions on how to change the network settings on the management station and how to verify that the address has changed.

- c. If your network is an IPv4 network, check the subnet mask to verify that it is set to 255.255.255.0, which is the default setting.
- d. From a command prompt, ping the A IP to make sure it is accessible.

Example

```
> ping 192.168.128.102
```

```
Reply from 192.168.128.102: bytes = 32 time<1ms TTL = 64
```

```
Ping statistics for 192.168.128.102:
```

```
Packets: Sent = 4, Received =4, Lost = 0 (0% loss)
```

```
Approximate round trip times in milli-seconds:
```

```
Minimum = 0ms, Maximum = 0ms, Average = 0 ms
```

2. Change the networking configuration.

The procedure you use depends on the model number of your storage array.

- For E2800 and E5700 storage arrays, see [Configuring a management port using System Manager](#) on page 31.
 - For E2700 and E5600 storage arrays, see [Configuring a management port using SANtricity Storage Manager](#) on page 33.
3. Disconnect the Ethernet cable from your management station, and reconnect the Ethernet cables from the controllers into your regular network.
 4. Complete the steps necessary to change the management station's IP address back to what it was initially.

Configuring management port using System Manager

The controller includes an Ethernet port used for system management. If necessary, you can change its transmission parameters and IP addresses.

About this task

During this procedure, you select port 1 and then determine the speed and port addressing method. Port 1 connects to the network where the management client can access the controller and System Manager.

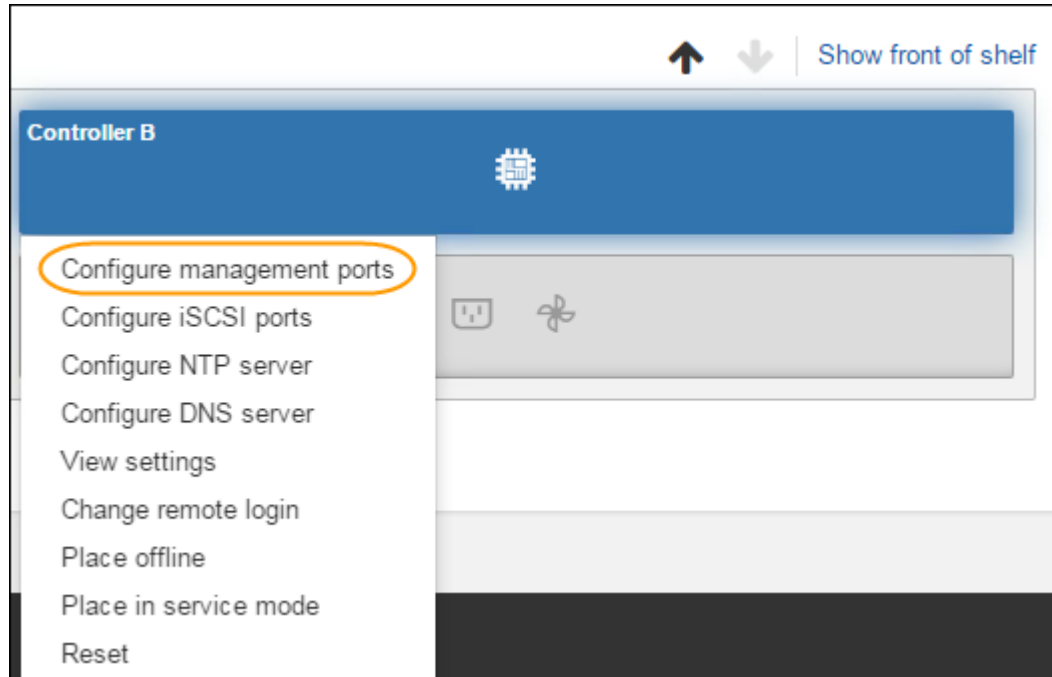
Note: Do not use port 2 on either controller. Port 2 is reserved for use by technical support.

Steps

1. Select **Hardware**.
2. If the graphic shows the drives, click **Show back of shelf**.
The graphic changes to show the controllers instead of the drives.
3. Click the controller with the management port you want to configure.

The controller's context menu appears.

4. Select **Configure management ports**.



The Configure Management Ports dialog box opens.

5. Make sure port 1 is displayed, and then click **Next**.
6. Select the configuration port settings, and then click **Next**.

Field Details

Field	Description
Speed and duplex mode	Keep the Auto-negotiate setting if you want System Manager to determine the transmission parameters between the storage array and the network; or if you know the speed and mode of your network, select the parameters from the drop-down list. Only the valid speed and duplex combinations appear in the list.
Enable IPv4 / Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks.

If you select **Enable IPv4**, a dialog box opens for selecting IPv4 settings after you click **Next**. If you select **Enable IPv6**, a dialog box opens for selecting IPv6 settings after you click **Next**. If you select both options, the dialog box for IPv4 settings opens first, and then after you click **Next**, the dialog box for IPv6 settings opens.

7. Configure the IPv4 and/or IPv6 settings, either automatically or manually.

Field Details

Field	Description
Automatically obtain configuration from DHCP server	Select this option to obtain the configuration automatically.

Field	Description
Manually specify static configuration	<p>Select this option, and then enter the controller's IP address. (If desired, you can cut and paste addresses into the fields.) For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.</p> <p>Attention: If you change the IP address configuration, you lose the management path to the storage array. Using the SANtricity Storage Manager Enterprise Management Window (EMW), you must remove the device from the EMW. Add it back to the EMW by selecting Edit > Add Storage Array, and then enter the new IP address. For more information, refer to the online help topics in the EMW.</p>

8. Click **Finish**.

Result

The management port configuration is displayed in the controller settings, Management Ports tab.

Configuring a management port using Storage Manager

Steps

1. Open the **SANtricity Storage Manager**.
The **Enterprise Management Window (EMW)** is displayed.
2. On the **Select Addition Method** screen, select the **Automatic** radio button, and then select **OK**.
This process finds all the storage arrays on the local sub-network. Several minutes might lapse to complete the process.
3. Name the storage array.
 - a. In the **EMW Setup** tab, select **Name/Rename Storage Arrays**.
 - b. In the **Select storage array** list, select the storage array you added.
 - c. In the **Storage array name** field, type a name for the storage array.
Storage array names must not exceed 30 characters and cannot contain spaces. Names can contain letters, numbers, underscores (_), hyphens(-), and pound signs (#). Choose a descriptive name for the storage array to make it easier for data center administrators to manage the storage resources over time.
 - d. Select **OK**.
4. Configure the network configuration information of the controllers, using information you obtain from your network administrator.
 - a. In the AMW, select the **Hardware** tab.
 - b. Select **Hardware > Controller > Configure > Management Ports**.
 - c. On the **Change Network Configuration** dialog box, select Controller A, Port 1 in the **Ethernet port** drop-down list.
 - d. From the **Speed and duplex mode** drop-down list, select **Auto-negotiate**.

Note: Attention Possible Connectivity Issues – After you select **Auto-negotiate**, make sure that your Ethernet switch also is set to **Auto-negotiate**.

- e. Depending on the format of your network configuration information, select the **Enable IPv4** check box, the **Enable IPv6** check box, or both check boxes.
- f. Depending on the format you have selected, enter the network configuration information (IP address, subnet mask, and gateway or IP address and routable IP address) in the **IPv4 Settings** tab or the **IPv6 Settings** tab.

Note: You must obtain the network configuration information from your network administrator.

- g. In the **Ethernet port** drop-down list, select Controller B, Port 1, and repeat step c through step f for controller B.
- h. Select **OK**.

Configuring multipath

If the express method for configuring the multipath driver does not meet the requirements of your configuration, you can consider alternate power methods.

Related concepts

[Power methods for configuring multipath](#) on page 42

[Configuring host utilities, virtualization, and clustering](#) on page 57

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Overview of multipath drivers

Multipath drivers help the hosts continue to operate without interruption when a physical path fails.

Multipath drivers provide a redundant path for the data cables connecting the storage array's controllers to the host bus adapters. For example, you can connect two host bus adapters to the redundant controller pair in a storage array, with different data cables for each controller. If one host bus adapter, one data cable, or one controller fails, the multipath driver automatically reroutes input/output (I/O) to the good path.

Multipath drivers provide these functions:

- They automatically identify redundant I/O paths.
- They automatically reroute I/O to an alternate controller when a controller fails or all of the data paths to a controller fail (failover).
- They check the state of known paths to the storage array.
- They provide status information on the controller and the bus.
- They check to see if Service mode is enabled on a controller and if the asymmetric logical unit access (ALUA) mode of operation has changed.
- They provide load balancing between available paths.

Multipath driver setup considerations

Most storage arrays contain two controllers that are set up as redundant controllers. If one controller fails, the other controller in the pair takes over the functions of the failed controller, and the storage array continues to process data. You can then replace the failed controller and resume normal operation. You do not need to shut down the storage array to perform this task.

The redundant controller feature is managed by the multipath driver software, which controls data flow to the controller pairs. This software tracks the current status of the connections and can perform the switch-over.

Whether your storage arrays have the redundant controller feature depends on a number of items:

- Whether the hardware supports it. Check to see whether you have duplex or simplex controllers in your configuration.

- Whether your OS supports certain multipath drivers. Refer to the installation and support guide for your operating system to determine whether your operating system supports redundant controllers.
- How the storage arrays are connected.

With the ALUA (I/O Shipping) feature, a storage array can service I/O requests through either controller in a duplex configuration; however, I/O shipping alone does not guarantee that I/O is routed to the optimized path.

Supported multipath drivers

E-Series storage arrays support multipath drivers specific to your operating system and a recommended host type.

This table provides general guidelines. Refer to the [Interoperability Matrix Tool](#) for compatibility information for specific HBA, multipath driver, OS level, and controller-drive tray support.

Operating System	Multipath driver	Recommended host type
Linux	DM-MP	If using the Automatic Load Balancing feature, use Linux DM-MP (Kernel 3.10 or later) along with the <code>scsi_dh_alua</code> device handler. Otherwise, use Linux DM-MP (Kernel 3.9 or earlier) along with the <code>scsi_dh_alua</code> device handler.
Linux	ATTO Multipath Director	Linux (ATTO) Note: You must use ATTO FC HBAs.
Linux	ATTO Multipath Director and clustered/parallel file system	ATTO Cluster/All OS Note: You must use ATTO FC HBAs and clustered/parallel file systems.
Linux	Symantec/Veritas Storage Foundation Multipath Driver	Linux (Symantec Storage Foundations)
Linux	VxDMP	Linux (Symantec Storage Foundation)

When you select either the **Typical (Full Installation)** option or the **Custom** installation option through the SMagent package, the host context agent is installed with SANtricity Storage Manager.

The preferred multipath driver is provided "in-box" with the operating system.

After the host context agent (SMagent) is installed, and the storage is attached to the host, the host context agent sends the host topology to the storage controllers through the I/O path. Based on the host topology, the storage controllers automatically define the host and the associated host ports, and set the host type. The host context agent sends the host topology to the storage controllers only once, and any subsequent changes made in SANtricity Storage Manager is persisted. For more information about where the host context agent resides in the install packaging, refer to [Storage management software components](#) on page 18.

If the host context agent does not select the recommended host type, you must manually set the host type in SANtricity software.

- To manually set the host type, from the Array Management Window, select the **Host Mappings** tab, select the host, and then select **Host Mappings >Host >Change Host Operating System**.
- If you are using SANtricity Storage Manager but not using partitions (for example, no Hosts defined), set the appropriate host type for the Default Group by selecting **Host Mappings > Default Group > Change Default Host Operating System**.
- If you are using SANtricity System Manager, use the "Create host manually" procedure in the System Storage Manager online help.

Related information

[NetApp Interoperability Matrix Tool](#)

Automatic Load Balancing feature overview

The Automatic Load Balancing feature provides automated I/O workload balancing and ensures that incoming I/O traffic from the hosts is dynamically managed and balanced across both controllers.

What is Automatic Load Balancing?

The Automatic Load Balancing feature provides improved I/O resource management by reacting dynamically to load changes over time and automatically adjusting volume controller ownership to correct any load imbalance issues when workloads shift across the controllers.

The workload of each controller is continually monitored and, with cooperation from the multipath drivers installed on the hosts, can be automatically brought into balance whenever necessary. When workload is automatically re-balanced across the controllers, the storage administrator is relieved of the burden of manually adjusting volume controller ownership to accommodate load changes on the storage array.

When Automatic Load Balancing is enabled, it performs the following functions:

- Automatically monitors and balances controller resource utilization.
- Automatically adjusts volume controller ownership when needed, thereby optimizing I/O bandwidth between the hosts and the storage array.

Host types that support the Automatic Load Balancing feature

Even though Automatic Load Balancing is enabled at the storage array level, the host type you select for a host or host cluster has a direct influence on how the feature operates. When balancing the storage array's workload across controllers, the Automatic Load Balancing feature attempts to move volumes that are accessible by both controllers and that are mapped only to a host or host cluster capable of supporting the Automatic Load Balancing feature. This behavior prevents a host from losing access to a volume due to the load balancing process; however, the presence of volumes mapped to hosts that do not support Automatic Load Balancing affects the storage array's ability to balance workload. For Automatic Load Balancing to balance the workload, the multipath driver must support TPGS and the host type must be included in the following table.

Host type supporting Automatic Load Balancing	With this multipath driver
Windows or Windows Clustered	MPIO with NetApp E-Series DSM
Linux DM-MP (Kernel 3.10 or later)	DM-MP with <code>scsi_dh_alua</code> device handler

Host type supporting Automatic Load Balancing	With this multipath driver
VMware	Native Multipathing Plugin (NMP) with VMW_SATP_ALUA Storage Array Type plug-in

Note: With minor exceptions, host types that do not support Automatic Load Balancing continue to operate normally whether or not the feature is enabled. One exception is that if a system has a failover, storage arrays move unmapped or unassigned volumes back to the owning controller when the data path returns. Any volumes that are mapped or assigned to non-Automatic Load Balancing hosts are not moved.

See the [Interoperability Matrix Tool](#) for compatibility information for specific multipath driver, OS level, and controller-drive tray support.

Note: For a host cluster to be considered capable of Automatic Load Balancing, all hosts in that group must be capable of supporting Automatic Load Balancing.

Multipath configuration diagrams

You can configure multipath in several ways. Each configuration has its own advantages and disadvantages.

This section describes the following configurations:


- Single-host configuration
- Direct connect and fabric connect configurations

This section also describes how the storage management software supports redundant controllers.

Single-Host configuration

In a single-host configuration, the host system contains two host bus adapters (HBAs), with a port on each HBA connected to different controllers in the storage array. The storage management software is installed on the host. The two connections are required for maximum failover support for redundant controllers.

Although you can have a single controller in a storage array or a host that has only one HBA port, you do not have complete failover data path protection with either of those configurations. The cable and the HBA become a single point of failure, and any data path failure could result in unpredictable effects on the host system. For the greatest level of I/O protection, provide each controller in a storage array with its own connection to a separate HBA in the host system.



1. Host System with Two SAS, Fibre Channel, iSCSI, or InfiniBand Host Bus Adapters
2. SAS, Fibre Channel, iSCSI, iSER over Infiniband or SRP over InfiniBand Connection – The Network Protocol Connection Might Contain One or More Switches
3. Storage Array with Two Controllers

Direct connect and fabric connect configurations


In a direct connect or fabric connect configuration, two host systems are each connected by two connections to both of the controllers in a storage array. SANtricity Storage Manager, including multipath driver support, is installed on each host.

Not every operating system supports this configuration. Consult the restrictions in the installation and support guide specific to your operating system for more information. Also, the host systems must be able to handle the multi-host configuration. Refer to the applicable hardware documentation.

In either a direct connect or fabric connect configuration, each host has visibility to both controllers, all data connections, and all configured volumes in a storage array.

The following conditions apply to these both direct connect and fabric connect configurations:


- Both hosts must have the same operating system version installed.
- The multipath driver configuration might require tuning.
- A host system might have a specified volume or volume group reserved, which means that only that host system can perform operations on the reserved volume or volume group.



1. Two Host Systems, Each with Two SAS, Fibre Channel, iSCSI, iSER over Infiniband, or SRP over InfiniBand Host Bus Adapters
2. SAS, Fibre Channel, iSCSI, iSER over Infiniband, or SRP over InfiniBand Connections with Two Switches (Might Contain Different Switch Configurations)
3. Storage Array with Two Controllers

Supporting redundant controllers

The following figure shows how multipath drivers provide redundancy when the host application generates a request for I/O to controller A, but controller A fails. Use the numbered information to trace the I/O data path.



1. Host Application
2. I/O Request
3. Multipath Driver
4. Host Bus Adapters
5. Controller A Failure
6. Controller B
7. Initial Request to the HBA
8. Initial Request to the Controller Failed
9. Request Returns to the Multipath Driver
10. Failover Occurs and I/O Transfers to Another Controller
11. I/O Request Re-sent to Controller B

How a multipath driver responds to a data path failure

One of the primary functions of the multipath driver is to provide path management. Multipath drivers monitor the data path for devices that are not working correctly or for multiple link errors.

If a multipath driver detects either of these conditions, the failover driver automatically performs the following steps:

- The multipath driver checks for the redundant controller.
- The multipath driver performs a path failure if alternate paths to the same controller are available. If all of the paths to a controller are marked offline, the multipath driver performs a controller failure. The failover driver provides notification of an error through the OS error log facility.
- For multipath drivers that are not using ALUA drivers, the multipath driver transfers volume ownership to the other controller and routes all I/O to the remaining active controller.
- For ALUA-based multipath drivers, controller B redirects I/O to the surviving controller (controller B). Then, if controller A is still active, controller B ships the I/O to controller A (SAN path loss case). If controller A has failed, controller B triggers a forced ownership transfer from the failed controller to itself (controller B).

User responses to a data path failure

You can use the Major Event Log (MEL) to troubleshoot a data path failure.

The information in the MEL provides the answers to these questions:

- What is the source of the error?
- What is required to fix the error, such as replacement parts or diagnostics?

When troubleshooting, follow these guidelines:

- Under most circumstances, contact technical support any time a path fails and the storage array notifies you of the failure.
- Use the MEL to diagnose and fix the problem, if possible.
- If your controller has failed and your storage array has customer-replaceable controllers, replace the failed controller. Follow the instructions provided with the controller.

Related information

[NetApp E-Series Systems Documentation Center](#)

Failover drivers for the Linux operating system

The Device Mapper Multipath (DM-MP) failover, which uses the Device Mapper generic framework for mapping one block device onto another, is supported with the Linux operating system. Device mapper is used for LVM, multipathing, and more.

Power methods for configuring multipath

Depending on your requirements, such as dividing I/O activity between RAID controllers or handling compatibility and migration, you can use the power methods for configuring multipath drivers.

Dividing I/O activity between two RAID controllers to obtain the best performance

For the best performance of a redundant controller system, use the storage management software to divide I/O activity between the two RAID controllers in the storage array. You can use a graphical user interface (GUI) or the command line interface (CLI).

The Automatic Load Balancing feature enables the system to dynamically reassign ownership so it can optimize the bandwidth between the hosts and the storage array. Note the following guidelines:

- If the Automatic Load Balancing feature is enabled, you do not need to perform the management tasks described in this section.
- If Automatic Load Balancing is enabled, you can select a preferred owner for a new volume when it is created, because there is no load history on that volume yet.
- By default, whenever possible the multipath driver directs I/O at the controller that is the preferred owner. This default method applies whether either of the following is true:
 - Preferred ownership is assigned automatically (Automatic Load Balancing is enabled).
 - Preferred ownership is assigned manually (Automatic Load Balancing is disabled).
- If you choose to disable Automatic Load Balancing, perform the management tasks described in this section to divide I/O activity between the two RAID controllers in the storage array.

To use the GUI to divide I/O activity between two RAID controllers, perform one of these procedures:

- From the SANtricity Storage Manager Array Management Window:
 - **Specify the owner of the preferred controller of an existing volume** – Select **Volume > Change > Ownership/Preferred Path**.

Note: You also can use this method to change the preferred path and ownership of all volumes in a volume group at the same time.
 - **Specify the owner of the preferred controller of a volume when you are creating the volume** – Select **Volume > Create**.

- From SANtricity System Manager:

Specify the owner of the preferred controller of an existing volume

1. Select **Storage > Volumes**.
2. Select any volume and then select **More > Change ownership**.

The **Change Volume Ownership** dialog box appears.

All volumes on the storage array appear in this dialog box.

3. Use the **Preferred Owner** drop-down list to change the preferred controller for each volume that you want to change, and confirm that you want to perform the operation.
- Using the CLI:
Go to the "Create RAID Volume (Free Extent Based Select)" online help topic for the command syntax and description.

Note: The volume might not use the new I/O path until the multipath driver reconfigures to recognize the new path. This action usually takes less than five minutes.

Installing DM-MP

All of the components required for DM-MP are included on the installation media.

Before you begin

You have installed the required packages on your system.

- For Red Hat (RHEL) hosts, verify the packages are installed by running `rpm -q device-mapper-multipath`.
- For SLES hosts, verify the packages are installed by running `rpm -q multipath-tools`.

About this task

By default, DM-MP is disabled in RHEL and SLES. Complete the following steps to enable DM-MP components on the host.

If you have not already installed the operating system, use the media supplied by your operating system vendor.

Steps

1. Use the procedures in the [Setting up the multipath.conf file](#) on page 50 to update and configure the `/etc/multipath.conf` file.
2. Do one of the following to enable the `multipathd` daemon on boot.

If you are using...	Do this ...
For RHEL 6.x systems:	<code>chkconfig multipathd on</code>
For SLES 11.x systems:	<ol style="list-style-type: none"> a. <code>chkconfig multipathd on</code> b. <code>chkconfig boot.multipath on</code>
For RHEL 7.x and SLES 12.x systems:	<code>systemctl enable multipathd</code>

3. Rebuild the `initramfs` image or the `initrd` image under `/boot` directory:

If you are using...	Do this ...
For SLES 11.x systems:	<pre>mkinitrd -k /boot/vmlinux-<flavour> -i /boot/ initrd-<flavour>.img -M /boot/System.map- <flavour></pre> <p>In this command, <flavour> is replaced with running kernel version from command "uname -r".</p>
For RHEL 6.x 7.x and 12.x systems:	<pre>dracut --force --add multipath</pre>

4. Make sure that the newly created /boot/initramfs-* image or /boot/initrd-* image is selected in the boot configuration file.

For example, for grub it is /boot/grub/menu.lst and for grub2 it is /boot/grub2/menu.cfg.
5. Do one of the following to verify and, if necessary, change the host type.
 - If you have hosts defined in the **SANtricity Storage Manager Host Mappings View**, go to step 6.
 - If you do not have hosts defined, right-click the default host group in the **SANtricity Storage Manager Host Mappings View**, and then set the default host type to either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. . Go to step 7.
 - If you are using SANtricity System Manager, use the "Create host manually" procedure in the online help to check the whether the hosts are defined and update as necessary in step 7.
6. In the SANtricity Storage Manager mappings view, right-click the host, and then select **Change Host Operating System**.
7. Verify that the selected host type is either **Linux DM-MP (Kernel 3.10 or later)** if you enable the Automatic Load Balancing feature, or **Linux DM-MP (Kernel 3.9 or earlier)** if you disable the Automatic Load Balancing feature. If necessary, change the selected host type to the appropriate setting.
8. Reboot the host.

Overview of migrating to the Linux DM-MP multipath driver

Because the MPP/RDAC driver is no longer available with SANtricity 11.25 or later, you must migrate to the Linux Device Mapper Multipath (DM-MP) driver. This procedure results in use of the Linux DM-MP (Kernel 3.9 or earlier) host type that uses the `scsi_dh_rdac` device handler. If you have already migrated to the DM-MP multipath driver in SANtricity 11.25 or later, you do not need to perform these steps.

If your system is not running the legacy MPP/RDAC driver, you do not need to complete this migration and can go to [Verifying OS compatibility with the Automatic Load Balancing feature](#) on page 49.

Migration consists of three steps: preparing for the migration, migrating the MPP/RDAC multipath driver to the Linux DM-MP driver, and verifying the migration to the Linux DM-MP driver.

Preparing for migration is non-disruptive and can be done ahead of time to ensure the system is ready for migration. Migrating to the Linux DM-MP Driver is disruptive because it involves a host reboot.

Downtime for the overall migration procedure involves time taken for the following actions and varies depending on different configurations and running applications:

- Application shutdown procedure
- Host Reboot procedure

Supported operating systems

Refer to the [NetApp Interoperability Matrix Tool](#) for supported OS versions for Device Mapper-multipath driver and storage array SANtricity OS (controller software) version. If your operating system and storage array SANtricity OS are not in the support matrix for the DM-MP driver, contact technical support.

Preparing to migrate to the DM-MP multipath driver

Because the MPP/RDAC driver is no longer available with SANtricity 11.25, you must migrate to the Linux Device Mapper Multipath (DM-MP) driver.

About this task

The system must be configured to use only persistent device names across all configuration files. This is suggested by all operating system vendors as well. These names are indicated by conventions like `/dev/disk/by-uuid` or `/dev/disk/by-label`. Persistent names are required because names like `/dev/sda` or `/dev/sdb` might change on the system reboot, depending on the SCSI device discovery order. Hard coded names can lead to devices disappearing and render the system unable to boot.

To configure persistent device naming conventions in your system, refer to your operating system vendor storage administration guide. NetApp has no recommendation about using specific conventions, provided that the chosen convention is verified by the user and supported by your operating system vendor.

For example, file system table configuration (`/etc/fstab`) should mount devices and partitions using either `/dev/disk/by-uuid` or `/dev/disk/by-label` symbolic names.

Steps

1. Mount devices by corresponding `/dev/disk/by-uuid` names instead of `/dev/sd` names:

```
UUID=88e584c0-04f4-43d2-ad33-ee9904a0ba32 /iomnt-test1 ext3 defaults
0 2
UUID=2d8e23fb-a330-498a-bae9-5df72e822d38 /iomnt-test2 ext2 defaults
0 2
UUID=43ac76fd-399d-4a40-bc06-9127523f5584 /iomnt-test3 xfs defaults 0
2
```

2. Mount devices by diskname labels:

```
LABEL=db_vol /iomnt-vg1-lvol ext3 defaults 0 2
LABEL=media_vol /iomnt-vg2-lvol xfs defaults 0 2
```

3. Make sure the boot loader configuration file (`/boot/grub/menu.lst` for `grub`) uses matching naming conventions.

For example, boot loader configurations using filesystem UUID or Label appear as the bold-faced labels in the following two examples:

```
linux /@/boot/vmlinuz-3.12.14-1-default root=UUID=e3ebb5b7-92e9-4928-aa33-55e2883b4c58
linux /@/boot/vmlinuz-3.12.14-1-default root=Label=root_vol
```

If you check for `SMdevices` after migration, you should see one device for each path.

Migrating the MPP/RDAC driver to the Linux DM-MP driver

Migrating from the MPP/RDAC multipath driver to the Linux DM-MP multipath driver allows you to ensure you have a supported multipath failover solution in SANtricity 11.30.

Steps

1. Uninstall the MPP/RDAC driver.

Typically the default location for the RDAC source directory is under the `/opt/StorageManager/` file path.

- If the MPP/RDAC driver is installed from the source, go to the RDAC source directory and then run the `#make uninstall` command.
- If the MPP/RDAC driver is installed from RPM, find the `linuxrdac` package name by specifying the `#rpm -q linuxrdac` command and then using the `#rpm -e "RDAC rpm name"` command to remove it from the system.

Note: Even after uninstalling the MPP/RDAC driver, make sure driver modules (`mppVhba.ko` and `mppUpper.ko`) remain loaded and running on the system so that application I/O is not disrupted. The host reboot performed in step 6 is necessary to unload these modules.

2. Using a text editor, replace the RDAC-generated initial ram disk image (`/boot/mpp-`uname -r`.img`) in the boot loader configuration file (for example, `/boot/grub/menu.lst` if using the GRUB boot loader) with the original RAM disk image from when you installed the operating system (that is, `/boot/initrd-<kernel version>.img` or `/boot/initramfs-<kernel version> file`).
3. Install and configure the Linux DM-MP multipath driver.
Refer to the [Installing DM-MP](#) on page 43 section to enable and configure the Linux in-box multipath driver. For supported OS versions for DM-MP driver, refer to the [Interoperability Matrix Tool](#).
4. Make sure you properly shut down all your applications.
5. Configure the HBA timeout values for the DM-MP driver, as recommended in the [NetApp Interoperability Matrix Tool](#). In some cases, DM-MP requires different values than MPP/RDAC, so make sure you verify these settings.
6. Reboot the host.
7. Verify that all file systems are mounted correctly by running the `mount` command.
 - If any of the file systems are not mounted check the `/etc/fstab` file for the corresponding device mount parameters provided.
 - If `/dev/sd` device names are used, change them to either `/dev/disk/by-uuid` symbolic link names or `/dev/mapper/` symbolic names.

Verifying the migration to Linux DM-MP driver

You can verify that the migration from the MPP/RDAC multipath driver to the Linux DM-MP multipath driver has been successful.

About this task

After both migration to DM-MP and the host reboot, `SMdevices` should show multiple entries for the same device, because it should display one device per path.

Steps

1. Verify that DM-MP device maps are created for all devices with NetApp/LSI vendor ID. Also verify that the path states are active and running .

The priority values for both priority groups of paths should be 14 and 9 respectively as shown in the following example. The hardware handler should be rdac and path selector should default as selected by operating system vendors.

```
# multipath -ll
```

```
mpatho (360080e50001b076d0000cd3251ef5eb0) dm-7 LSI ,INF-01-00
size=5.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle'
hwhandler='1 rdac' wp=rw
|+- policy='service-time 0' prio=14 status=active
| |- 5:0:1:15 sdag 66:0 active ready running
| - 6:0:1:15 sdbm 68:0 active ready running
`-+ policy='service-time 0' prio=9 status=enabled
| - 5:0:0:15 sdq 65:0 active ready running
- 6:0:0:15 sdaw 67:0 active ready running
```

```
# multipathd show paths
```

```
hcil dev dev_t pri dm_st chk_st dev_st next_check
5:0:0:0 sdb 8:16 14 active ready running XXXXXXXX... 14/20
5:0:0:1 sdc 8:32 9 active ready running XXXXXXXX... 14/20
5:0:0:10 sdl 8:176 9 active ready running XXXXXXXX... 14/20
5:0:0:11 sdm 8:192 14 active ready running XXXXXXXX... 14/20
```

```
# multipathd show maps
```

```
name sysfs uuid
mpathaa dm-0 360080e50001b081000001b525362ff07
mpathj dm-1 360080e50001b076d0000cd1a51ef5e6e
mpathn dm-2 360080e50001b076d0000cd2c51ef5e9f
mpathu dm-3 360080e50001b08100000044a51ef5e2b
```

If any of the path states appear as "ghost," make sure that Linux DM-MP (Kernel 3.9 or earlier) host type is selected from SANtricity Storage Manager Host Mapping view. If any path states appear as "faulty" or "failed" refer to [Troubleshooting Device Mapper](#) on page 56. If you require further assistance, contact technical support.

If none of the NetApp/LSI devices appear with these commands, check the `/etc/multipath.conf` file to see if they are blacklisted. If so, remove those blacklisted entries, and then rebuild the initial RAM disk as mentioned in step 2 of [Migrating the MPP/RDAC driver to the Linux DM-MP driver](#) on page 46.

2. If LVM is configured, run the following commands, and then verify that all the VG/LV/PV devices are referenced by either WWID or "mpath" names rather than `/dev/sd` device names.

```
# pvdisplay
```

```
--- Physical volume ---
```

```
PV Name          /dev/mapper/mpathx_part1
VG Name          mpp_vg2
PV Size          5.00 GiB / not usable 3.00 MiB
Allocatable      yes
PE Size          4.00 MiB
Total PE         1279
Free PE          1023
```



```
Allocated PE    256
PV UUID        v671wB-xgFG-CU0A-yjc8-snCc-d29R-ceR634
```

```
# vgdisplay
```

```
--- Volume group ---
```

```
VG Name          mpp_vg2
System ID
Format           lvm2
Metadata Areas   2
Metadata Sequence No 2
VG Access        read/write
VG Status        resizable
MAX LV           0
Cur LV          1
Open LV          1
Max PV           0
Cur PV          2
Act PV           2
VG Size          9.99 GiB
PE Size          4.00 MiB
Total PE         2558
Alloc PE / Size  512 / 2.00 GiB
Free PE / Size   2046 / 7.99 GiB
VG UUID          jk2xgS-9vS8-ZMmk-EQdT-TQRi-ZUNO-RDgPJz
```

```
# lvdisplay
```

```
--- Logical volume ---
```

```
LV Name          /dev/mpp_vg2/lvol0
VG Name          mpp_vg2
LV UUID          tFGMy9-eJhk-FGxT-XvbC-ItKp-BGnI-bzA9pR
LV Write Access  read/write
LV Creation host, time a7-boulevard, 2014-05-02 14:56:27 -0400
LV Status        available
# open           1
LV Size          2.00 GiB
Current LE       512
Segments         1
Allocation       inherit
Read ahead sectors auto
- currently set to 1024
Block device     253:24
```

3. If you encounter any issues, perform the appropriate file system checks on the devices.

Verifying correct operational mode for ALUA

After setting up the DM-MP multipath driver, make sure that the DM-MP configuration is set up correctly and that ALUA mode is operational.

Steps

1. Type `SMdevices`.

If operating in the correct mode, the output displays either `Active/Optimized` or `Active/Non-optimized` at the end of each line. The host can see the LUNs that are mapped to it.

If not operating in the correct mode, the output displays as either `passive` or `unowned`.

2. At the command prompt, type `multipath -ll`.

If both controllers are online, there should be exactly two path groups for each LUN, one for each controller.

If your system is running Automatic Load Balancing, the output should look as follows. Note the priority values of 50 and 10, and `hwhandler='1 alua'`. These indicate the correct mode for DM-MP with Automatic Load Balancing:

```
360080e50001b076d0000cd2451ef5e8a dm-8 NETAPP,INF-01-00
size=5.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-- policy='round-robin 0' prio=50 status=enabled
|  |-- 8:0:2:2 sdn 8:208 active ready running
|  |-- 8:0:3:2 sdy 65:128 active ready running
|-- policy='round-robin 0' prio=10 status=enabled
|  |-- 7:0:2:2 sdc 8:32 active ready running
|  |-- 7:0:3:2 sdaj 66:48 active ready running
```

If your system is *not* running Automatic Load Balancing, the output should look as follows. Look for priority values greater than 8 and `hwhandler='1 rdac'`. These indicate the correct operational mode for DM-MP without Automatic Load Balancing:

```
360080e50001b076d0000cd2451ef5e8a dm-8 NETAPP,INF-01-00
size=5.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 rdac' wp=rw
|-- policy='round-robin 0' prio=14 status=enabled
|  |-- 8:0:2:2 sdn 8:208 active ready running
|  |-- 8:0:3:2 sdy 65:128 active ready running
|-- policy='round-robin 0' prio=9 status=enabled
|  |-- 7:0:2:2 sdc 8:32 active ready running
|  |-- 7:0:3:2 sdaj 66:48 active ready running
```

Verifying OS compatibility with the Automatic Load Balancing feature

Verify OS compatibility with the Automatic Load Balancing feature before setting up a new (or migrating an existing) system.

About this task

If you would like to enable Automatic Load Balancing with your RHEL 7 or SLES 12 Linux system, please complete the following steps to verify compatibility with the specific operating system before setting up a new (or migrating an existing) system..

Steps

1. Go to the [Interoperability Matrix Tool](#) to find your solution and verify support.
If your system is running Red Hat Enterprise Linux 6 or SUSE Linux Enterprise Server 11, please refer to [Knowledge base article 1015843](#).
2. Use the procedures in [Setting up the multipath.conf file](#) on page 50 to update and configure the `/etc/multipath.conf` file.
3. Ensure that both `retain_attached_device_handler` and `detect_prio` are set to `yes` for vendor "NETAPP" and product "INF-01-00", or use default settings.

Migrating an existing DM-MP system to use Automatic Load Balancing

Make sure your SANtricity 11.30 system is prepared to use the Automatic Load Balancing feature to balance workloads across your configuration.

Before you begin

You have completed the steps in [Verifying OS Compatibility with Auto-Load Balancing](#) on page 49.

About this task

Enable Auto-Load Balancing support for existing DM-MP systems.

Steps

1. In the SANtricity Storage Manager mappings view, right-click the host, and then select **Change Host Operating System**.
2. Verify the selected host type is `Linux DM-MP (Kernel 3.10 or later)`. If necessary, change the selected host type to the Host type name.
3. Reboot the host.

Setting up the multipath.conf file

The `multipath.conf` file is the configuration file for the multipath daemon, `multipathd`. The `multipath.conf` file overrides the built-in configuration table for `multipathd`. Any line in the file whose first non-white-space character is `#` is considered a comment line. Empty lines are ignored.

Note: For SANtricity operating system 8.30, NetApp recommends using the default settings as provided.

Example `multipath.conf` are available in the following locations:

- For SLES, `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic`
- For RHEL, `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf`

All the lines in the sample `multipath.conf` file are commented out. The file is divided into five sections:

- **defaults** – Specifies all default values.
- **blacklist** – All devices are blacklisted for new installations. The default blacklist is listed in the commented-out section of the `/etc/multipath.conf` file. Blacklist the device mapper multipath by WWID if you do not want to use this functionality.
- **blacklist_exceptions** – Specifies any exceptions to the items specified in the section `blacklist`.
- **devices** – Lists all multipath devices with their matching vendor and product values.
- **multipaths** – Lists the multipath device with their matching WWID values.

The DM-MP multipath driver has built-in default values, as well as built-in settings for different vendor and product ID combinations. When defining sections in `multipath.conf`, it has the following effects:

- Parameter values defined in the `defaults` section merge with the built-in defaults, replacing those values.

- Parameter values defined in a device section merge with the built-in defaults for that vendor and product ID if the device already exists in the built-in configuration. To ensure this merging occurs, the vendor and product strings must match the built-in configuration exactly.
- For each parameter, the value is determined in the following sequence:
 1. If defined, the multipath section for each device.
 2. If defined, the device section for the device's vendor and product ID.
 3. The internal default value.

In the following tasks, you modify the default, blacklist and devices sections of the `multipath.conf` file. Remove the initial `#` character from the start of each line you modify.

Updating the blacklist section

With the default settings, UTM LUNs might be presented to the host. I/Os operations, however, are not supported on UTM LUNs. To prevent I/O operations on the UTM LUNs, add the vendor and product information for each UTM LUN to the blacklist section of the `/etc/multipath.conf` file.

About this task

This optional task allows you to remove virtualized disks, as well as disks from other storage arrays that are managed by other multipath drivers.

Note: For SANtricity 8.40, NetApp recommends that you verify the default settings and do not modify any settings.

The entries should follow the pattern of the following example.

```
blacklist {
    device {
        vendor "*"
        product "Universal Xport"
    }
}
```

Updating the devices section of the multipath.conf file

If your host is running RHEL 6.5, SLES 11.3 or an earlier version, you can update the `/etc/multipath.conf` file. If you are using a later release, simply create an empty `/etc/multipath.conf` file so the system automatically applies all the default configurations, which includes supported values for NetApp E-Series and EF-Series devices.

The following example shows part of the `devices` section in the `/etc/multipath.conf` file. The example shows the vendor ID as `NETAPP` or `LSI` and the product ID as `INF-01-00`. Modify the `devices` section with product and vendor information to match the configuration of your storage array. If your storage array contains devices from more than one vendor, add additional `device` blocks with the appropriate attributes and values under the `devices` section. NetApp has no recommendation on a particular path selector to use. Therefore, the default path selector is selected with the device settings as shown in the example. The command `"multipathd show config"` shows the path selector in the defaults section.

Note: Only update the devices section of the `/etc/multipath.conf` file if your host is running RHEL 6.5, SLES 11.3, or an earlier version. For Cluster configurations, set failback to manual as specified in the [Interoperability Matrix Tool](#).

```
devices {
    device {
        vendor          "(LSI|NETAPP)"
        product         "INF-01-00"
```

```

path_grouping_policy      group_by_prio
detect_prio               yes
prio                     rdac
path_checker              rdac
hardware_handler          "1 rdac"
failback                  immediate
features                  "2 pg_init_retries 50"
no_path_retry             30
retain_attached_hw_handler yes
    }
}

```

Note: Other than the parameters listed above, the internal default value for all other parameters is the recommended value and should not be changed. Both the round robin or service time path selectors are fully supported, and NetApp makes no specific recommendation for either one.

Attribute	Parameter value	Description
path_grouping_policy	group_by_prio	The path grouping policy to be applied to this specific vendor and product storage.
detect_prio	yes	The system detects the path policy routine.
prio	rdac	The program and arguments to determine the path priority routine. The specified routine should return a numeric value specifying the relative priority of this path. Higher numbers have a higher priority.
path_checker	rdac	The method used to determine the state of the path.
hardware_handler	"1 rdac"	The hardware handler to use for handling device-specific knowledge.
failback	immediate	A parameter to tell the daemon how to manage path group failback. In this example, the parameter is set to <code>immediate</code> , so failback occurs in the next check interval in which the path is available. To disable the failback, set this parameter to <code>manual</code> . When clustering or shared LUN environments are used without the Automatic Load Balancing feature enabled, set this parameter to <code>manual</code> .
features	"2 pg_init_retries 50"	Features to be enabled. This parameter sets the kernel parameter <code>pg_init_retries</code> to 50. The <code>pg_init_retries</code> parameter is used to retry the mode select commands.

Attribute	Parameter value	Description
<code>no_path_retry</code>	30	Specify the number of retries before queuing is disabled. Set this parameter to <code>fail</code> for immediate failure (no queuing). When this parameter is set to <code>queue</code> , queuing continues indefinitely. The amount of time is equal to the parameter value multiplied by the <code>polling_interval</code> (usually 5), for example, 150 seconds for a <code>no_path_retry</code> value of 30.
<code>retain_attached_hw_handler</code>	yes	Specifies that the current hardware handler continues to be used.

Setting up DM-MP for large I/O blocks

When a single I/O operation requests a block larger than 512 KB, this is considered to be a large block. You must tune certain parameters for a device that uses Device Mapper Multipath (DM-MP) for the device to perform correctly with large I/O blocks. Parameters are usually defined in terms of blocks in the kernel, and are shown in terms of kilobytes to the user. For a normal block size of 512 bytes, simply divide the number of blocks by 2 to get the value in kilobytes.

About this task

The following parameters affect performance with large I/O blocks:

- `max_hw_sectors_kb` (RO) - This parameter sets the maximum number of kilobytes that the hardware allows for request.
- `max_sectors_kb` (RW) - This parameter sets the maximum number of kilobytes that the block layer allows for a file system request. The value of this parameter must be less than or equal to the maximum size allowed by the hardware. The kernel also places an upper bound on this value with the `BLK_DEF_MAX_SECTORS` macro. This value varies from distribution to distribution.
- `max_segments` (RO) - This parameter enables the low-level driver to set an upper limit on the number of hardware data segments in a request. In the HBA drivers, this is also known as `sg_tablesize`.
- `max_segment_size` (RO) - This parameter enables the low-level driver to set an upper limit on the size of each data segment in an I/O request in bytes. If clustering is enabled on the low-level driver, it is set to 65536 or it is set to system `PAGE_SIZE` by default, which is typically 4K. The maximum I/O size is determined by the following:

```
MAX_IO_SIZE_KB = MIN(max_sectors_kb, (max_segment_size *
max_segments) / 1024)
```

Steps

1. Set the value of the `max_segments` parameter for the respective HBA driver as load a time module parameter.

The following table lists HBA drivers which provide module parameters to set the value for `max_segments`.

HBA	Module parameter
LSI SAS (mpt2sas)	max_sgl_entries
Emulex (lpfc)	lpfc_sg_seg_cnt
InfiniBand (ib_srp)	cmd_sg_entries
Brocade (bfa)	bfa_io_max_sge

- If supported by the HBA, set the value of `max_hw_sectors_kb` for the respective HBA driver as a load time module parameter.

This parameter is in sectors and is converted to kilobytes.

HBA	Parameter	How to set
LSI SAS (mpt2sas)	<code>max_sectors</code>	Module parameter
Infiniband (ib_srp)	<code>max_sect</code>	Open <code>/etc/srp_daemon.conf</code> and add <code>"a max_sect=<value>"</code>
Brocade (bfa)	<code>max_xfer_size</code>	Module parameter

- Set the value for the `max_sectors_kb` parameter for all physical paths for dm device in sysfs. For example, `echo N >/sys/block/sd device name /queue/max_sectors_kb`. In the command, *N* is an unsigned number less than the `max_hw_sectors_kb` value for the device; *sd device name* is the name of the sd device.
- On the command line, enter the command `echo N >/sys/block/dm device name /queue/max_sectors_kb` to set the value for the `max_sectors_kb` parameter for all dm device in sysfs. In the command, *N* is an unsigned number less than the `max_hw_sectors_kb` value for the device; *dm device name* is the name of the dm device represented by dm-X.

Using the device mapper devices

Multipath devices are created under `/dev/` directory with the prefix `dm-`. These devices are the same as any other block devices on the host. To list all of the multipath devices, run the `multipath -ll` command.

The following example shows system output from the `multipath -ll` command for one of the multipath devices.

```
360080e50001b076d0000cd8451ef5fea dm-0 NETAPP,INF-01-00
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|  |- 4:0:1:0 sdc 8:32 active ready running
|  `-- 4:0:2:0 sdd 8:48 active ready running
`+- policy='service-time 0' prio=10 status=enabled
|  |- 1:0:2:0 sda 8:0 active ready running
|  `-- 1:0:3:0 sdb 8:16 active ready running
```

In this example, the multipath device nodes for this device are `/dev/mapper/360080e50001b076d0000cd8451ef5fea` and `/dev/dm-0`. This example shows how the output should appear during normal operation. The lines beginning with `"policy="` are the path groups. There should be one path group for each controller. The path group currently being used for I/O access will have a status of `active`.

The following table lists some basic options and parameters for the `multipath` command.

Command	Description
<code>multipath -h</code>	Prints usage information
<code>multipath</code>	With no arguments, attempts to create multipath devices from disks not currently assigned to multipath devices
<code>multipath -ll</code>	Shows the current multipath topology from all available information, such as the sysfs, the device mapper, and path checkers
<code>multipath -ll map</code>	Shows the current multipath topology from all available information, such as the sysfs, the device mapper, and path checkers
<code>multipath -f map</code>	Flushes the multipath device map specified by the map option, if the map is unused
<code>multipath -F</code>	Flushes all unused multipath device maps

How to use partitions on DM devices

Multipath devices can be partitioned like any other block device. When you create a partition on a multipath device, device nodes are created for each partition. The partitions for each multipath device have a different dm- number than the raw device.

For example, if you have a multipath device with the WWID `3600a0b80005ab177000017544a8d6b9c` and the user friendly name `mpathb`, you can reference the entire disk through the following path:

```
/dev/mapper/mpathb
```

If you create two partitions on the disk, they are accessible through the following path:

```
/dev/mapper/mpathbp1
/dev/mapper/mpathbp2
```

If you do not have user friendly names enabled, the entire disk are accessible through the following path:

```
/dev/mapper/3600a0b80005ab177000017544a8d6b9c
```

And the two partitions are accessible through the following path:

```
/dev/mapper/3600a0b80005ab177000017544a8d6b9cp1
/dev/mapper/3600a0b80005ab177000017544a8d6b9cp2
```

Rescanning devices with the DM-MP multipath driver

In a Linux operating system, you can rescan SCSI devices to work with your new multipath DM-MP driver after installing the new driver.

You can rescan devices through the `rescan-scsi-bus.sh` script, which is included in the `sg3_utils` package.

Note: To use this functionality, the operating system requires a minimum version of `rescan-scsi-bus.sh`. All OSs supported by this release contain this functionality. The minimum package version is 1.27.

- For RHEL6, use the `sg3_utils-1.28-6.el6` package.

- For SLES 11, use the **sg3_utils-1.35-0.13.1** package.

Command	Description
<code>rescan-scsi-bus.sh -m</code>	Scan for new SCSI devices, and then attempt to create multipath devices.
<code>rescan-scsi-bus.sh -m -u</code>	Search for remapped SCSI devices, remove the old multipath device, and then create a new multipath device.
<code>rescan-scsi-bus.sh -m -s</code>	Search for resized SCSI devices, and then update the multipath device size.
<code>rescan-scsi-bus.sh -m -r</code>	Search for unmapped SCSI devices, and then delete both the devices and their associated multipath device.

Warning: Using the `-r` option in the `sg3_utils` packages for SLES 11 SP4 and SLES 12 SP 1 incorrectly deletes all disks, requiring you to reboot the hosts. . For SLES 11 SP4, upgrade to **sg3_utils-1.40-0.21.1**. For SLES 12 SP 1, a fix is still pending. Otherwise, on these OS versions, you should reboot the hosts to remove volumes that have been unmapped.

Troubleshooting Device Mapper

Situation	Resolution
Is the multipath daemon, <code>multipathd</code> , running?	At the command prompt, enter the command: <code>#service multipathd status</code> .
Why are no devices listed when you run the <code>multipath -ll</code> command?	At the command prompt, enter the command: <code>#cat /proc/scsi/scsi</code> . The system output displays all of the devices that are already discovered. Verify that the <code>multipath.conf</code> file has been updated with proper settings. You can check the running configuration with the <code>multipathd show config</code> command.

Configuring host utilities, virtualization, and clustering

For load balancing, availability, and security concerns, virtualization and clustering are essential considerations for your storage configuration. The Unified Host Utilities package provides tools to optimize NetApp storage performance, set required parameters on hosts, connect to hosts, and display information about storage cluster nodes.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Host utilities

The Host Utilities package provides the `sanlun` utility, which provides an end-to-end view of a LUN, as well as a limited view of the attached host bus adapters. This utility can be installed anytime on the server attached to storage and it does not alter any settings on the Linux operating system.

The `sanlun` command is the only command with the ability to show the SANtricity volume name mapping (to device name), as well as the current LUN ownership. Use the `-p` flag.

```
# sanlun lun show
controller(7mode/E-Series)/
vserver          lun name    device      host      lun
(cDOT/FlashRay) filename  adapter  protocol size  product
-----
-----
ICTM1922S01C3    H5-50      /dev/sdau   host26    iSCSI    2g    E-
Series
ICTM1922S01C2    H5-51      /dev/sdw    host9      FCP      2g    E-
Series
```

The following example shows the output with Auto-Load Balancing enabled:

```
# sanlun lun show -p
E-Series Array: a7-Tahoe-FC(60080e50001b0810000000004cbf6fe3)
Volume Name: V106
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
Mode: TPGS (Active/Active)
Preferred Path Auto Changeable: No (User Disabled)
Implicit Failback: Enabled
UTM LUN: Enabled (7)
LUN: 0
LUN Size: 2g
Product: E-Series
Host Device: 360080e50001b076d0000cd8451ef5fea
Multipath Policy: service-time 0
Multipath Provider: Native
-----
-----
host      controller
path      path      /dev/     host      controller
state     type      node     adapter   target
-----
-----
up        primary   sdc      host4     B3
up        primary   sdd      host4     B4
```

```
up          secondary sda    host1     A4
up          secondary sdb    host1     A3
```

The following example shows the output without Auto-Load Balancing enabled:

```
# sanlun lun show -p
root@ICTM1618S01H1:/opt/wendy# sanlun lun show -p

E-Series Array: ictm1618s01c1(60080e500029059c0000000055bf570d)
  Volume Name: NC-H1-127
Preferred Owner: Controller in Slot B
Current Owner: Controller in Slot B
  Mode: RDAC (Active/Active)
  UTM LUN: None
    LUN: 127
  LUN Size: 18g
  Product: E-Series
Multipath Provider: Unknown
-----
-----
host      controller
path      path      /dev/    host      controller
state     type      node     adapter   target
-----
-----
up        secondary sddj     host8     A1
up        primary   sdcg     host8     B1
```

Virtualization considerations

For the purpose of storage, virtualization refers to the act of creating a virtual machine (VM) within a parent operating system. Virtualization isolates applications, and allows for virtual desktop deployments that can provide security not available on the physical operating system. In addition, virtualization can ensure high availability while reducing hardware costs across an enterprise. There are many virtualization technologies built onto operating systems, as well as operating systems whose main purpose is to provide virtualization.

Virtualization offers a wide range of capabilities to an organization:

- **Server consolidation:** Many servers can be replaced by one large physical server, so hardware is consolidated, and guest operating systems are converted to virtual machines. This consolidation provides the ability to run legacy software on new hardware.
- **Isolation:** A guest operating system can be fully isolated from the host running it. If the virtual machine is corrupted, the host system is not harmed.
- **Migration:** A process to move a running virtual machine to another physical machine. Live migration is an extended feature that allows this move without disconnection of the client or the application.
- **Disaster recovery:** Virtualized guest systems are less dependent on the hardware.

For virtualization deployments on NetApp E-Series products, storage volume layout and host mappings should be considered. Additionally, host multipathing and connection Pass-Thru might be required.

Storage volume layout

When planning your volume layout, the following general guidelines apply:

- The larger the deployment, the higher the disk count.

If volume groups or disk pools are not large enough, latency problems can cause a series of timeouts.

- As the volumes used by virtual machines increases within a volume group, the IO workload moves from mostly sequential to mostly random in pattern.
For example, one VMs workload will look sequential, but if you provide a series of VMs, the expanded workload will look random over time.

Volume Mapping & Pass Through

Volumes are typically mapped to the parent directory. Unless there are multiple RAID groups, NetApp recommends using one large disk for VMs. The large disk can later be divided into smaller segments for virtualization.

If copy services backup individual VMs, then volumes need to be mapped for each VM to the parent operating system. Some virtual environments allow storage to be managed by the virtual machine directly. This management requires you to define an additional host and host-type on the storage array to be configured.

Volumes mapped to this host are not visible to the parent operating system.

Multipathing and virtualization

Virtualization must account for multipathing software. In a typical virtualized environment, the parent operating system performs any failover scenarios required. If the VM is a pass thru, any pathing considerations need to be handled through failover within the VM.

Virtualization needs to account for multipathing software. In a typical virtualized environment, the parent os performs any failover scenarios required. If the VM is a pass thru, any pathing considerations need to be handled through failover within the VM.

When planning your installation, consider the following methods:

- **Single Root I/O Virtualization (SR-IOV)** is a specification that allows a single Peripheral Component Interconnect Express (PCIe) physical device under a single root port to appear to be multiple separate physical devices to the hypervisor or the guest operating system.
- **N_Port ID Virtualization (NPIV)** is a Fibre Channel feature that allows multiple Fibre Channel node port (**N_Port**) IDs to share a single physical **N_Port**. Multiple Fibre Channel initiators can occupy a single physical port so each virtual server can see its own storage and no other virtual server's storage.

Linux-specific virtualization options

Depending on your operating system, there are virtualization guidelines that different operating systems provide.

Red Hat provides guidelines for setting up virtualization with Red Hat Enterprise Linux systems.

[Red Hat: Virtualization](#)

Host clustering support

Host clustering provides a way to load balance and make highly available applications. Generally, a cluster solution is one or more servers that work together and can be viewed as a single system. Cluster solutions improve performance and availability over a single computer, while being more cost-effective.

The following terms are common to a discussion of Host clustering:

Nodes

The underlying clients running the cluster application that make up the cluster. Traditionally, nodes pertained to physical servers, but some clustering packages allow virtual machines to also play the role of a node. In most cases, all nodes in a cluster use the same hardware and the same operating system.

Services

An entity shared by cluster nodes, Services are the high-level, tangible entities that depend on everything below them in the clustering hierarchy. Network shares and applications are examples of Services.

Services are monitored for accessibility and stability by the cluster application.

Resources

An entity shared by cluster nodes, Resources are a lower-level entity than Services. Resources include entities like disks, and IP addresses.

Resources are exposed through services and monitored for accessibility and stability by the cluster application.

Cluster accessibility

Managing accessibility is critical for all cluster nodes. The best methods for managing accessibility involve using a "heartbeat" for node-to-node communication, using "fencing" to control access to a cluster, and using a "quorum" to control the size of a cluster.

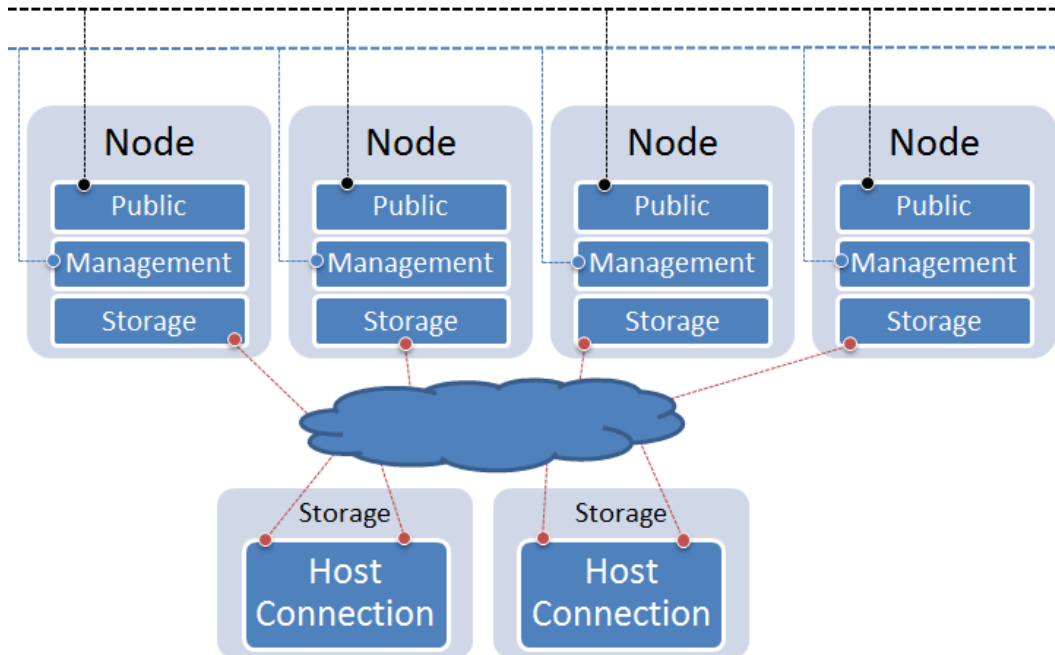
- **heartbeat:** All cluster nodes communicate with each other through a heartbeat. The most obvious communication method is through the network. If possible, the heartbeat should be on a separate network. Clusters can also use serial cables or shared disks for communications. The heartbeat is so vital, that in some clusters a single dropped packet can result in a fenced node.
- **fencing:** The process nodes use to kick other nodes from a cluster. This process varies among cluster packages and can happen for a variety of reasons. Clusters usually have multiple types of fencing devices (ways to remove nodes from a cluster) including APC Power and SCSI Reservations.
- **quorum:** Some clusters adopt the idea of a quorum: a cluster is not be established until enough nodes have joined and agree that a cluster can be started. If enough nodes leave and there is no longer a quorum, the cluster can dissolve. Quorums can be established from the network or from shared disks (where a disk is called the quorum disk). Normally, quorum disks are more tolerant to node failures as network quorum requires a node majority ($N/2+1$).

Most clusters also have the concept of a failover domain. The failover domain determines which node will own the service at which time and can usually prioritize service migrations for load balancing.

Other clusters claim a "master node" in cases of failure. This method is not widely used because if the master node fails, the cluster can become 'split brain'. Split brain occurs when nodes are still claimed as active but do not have communications to other nodes who also claim to be active. The consequences can be devastating as similar services acting on the same resource can overwrite one another.

Cluster topology

Cluster connections consist of a public network, a private, cluster management network, and a storage network.



- **Public Network:** this network provides access to the outside world or LAN.
- **Private Network:** It is recommended to isolate a network specifically for cluster management. Some clustering software allow different management types (serial, network, etc).
- **Storage Network:** Traditional connections to storage. This can be a variety of protocols.

Cluster shared storage in SANtricity

Allowing multiple hosts to share the same storage is critical in many clusters.

About this task

Shared storage can be used in couple of ways by the cluster.

- **Shared Disk File System:** Some file systems are distributed aware. These file systems typically deploy a rigorous concurrency model to keep incoming data requests serialized.
- **General Parallel File System (GPFS):** A high-performance clustered file system that can be deployed in either shared-disk or shared-nothing distributed parallel modes. GPFS provides higher I/O performance by striping blocks of data from individual files over multiple disks, and reading and writing these blocks in parallel.
- **Quorum Disk:** Shared storage can provide a disk to the cluster designed to keep the cluster operational. If a node cannot access the quorum disk, then the node understands that it is no longer part of the cluster until access become available. Nodes communicate through the quorum disk to relay state information. This disk can be used in place of a heartbeat and can be the trigger for fencing behavior within the cluster.

To create shared storage in both SANtricity Storage Manager and SANtricity System Manager (if your storage array has an E2800 controller shelf), use the following general procedure, supplemented with online help topics:

Steps

1. Create all of the individual hosts that will share access to a set of volumes.
2. Do one of the following:
 - If you have an E2700, E5600, or EF560 controller shelf, create a host group.
 - If you have an E2800 controller shelf, create a host cluster.
3. Add all of the individual hosts to the host cluster or the host group.
4. Map all volumes into the host group or assign all volumes to the host cluster that you want to share.

When complete, all hosts can see the volume.

What are SCSI reservations?

SCSI reservations allow a node to lock volume access to other nodes. There are two types in use: SCSI-2 reservations and SCSI-3 persistent reservations.

- SCSI-2 reservations provide two commands: `SCSI Reserve` and `SCSI Release`. A bus reset clears the LUN reservation. SCSI-2 reservations have been deprecated in recent standards, but are still available on various clusters.
- SCSI-3 persistent reservations, as its name suggests, provides reservation persistence across various resets. Exclusive LUN access is provided by registering, reserving, and locking the LUN. When a node wishes to relinquish the lock, the node releases the LUN. Additional registrations are not required to perform additional LUN reservations.

Deciding whether to use disk pools or volume groups

You can create volumes using either a disk pool or a volume group. The best selection depends primarily on your key storage requirements, such as expected I/O workload, performance requirements, and data protection requirements.

If you have a highly sequential workload and need maximum system bandwidth and the ability to tune storage settings, choose a volume group.

If you have a highly random workload and need faster drive rebuilds, simplified storage administration, and thin provisioning, choose a Dynamic Disk Pool (DDP).

Use case	Volume group	Dynamic Disk Pool
Workload - random	Good	Better
Workload - sequential	Better	Good
Drive rebuild times	Slower	Faster
Performance (optimal mode)	Good Best for large-block, sequential workloads	Good Best for small-block, random workloads
Performance (drive rebuild mode)	Degraded. Up to 40% drop in performance	Better
Multiple drive failure	Less data protection Slow rebuilds, greater risk of data loss	Greater data protection Faster, prioritized rebuilds
Adding drives	Slower Requires Dynamic Capacity Expansion operation	Faster Add to disk pool on the fly
Thin provisioning support	No	Yes
SSDs	Yes	Yes
Simplified administration	No Allocate global hot spares, configure RAID	Yes No hot spare or RAID settings to configure
Tunable performance	Yes	No

Creating a volume group

You use a volume group to create one or more volumes that are accessible to the host. A volume group is a container for volumes with shared characteristics such as RAID level and capacity.

About this task

With larger capacity drives and the ability to distribute volumes across controllers, creating more than one volume per volume group is a good way to make use of your storage capacity and to protect your data.

Follow these guidelines when you create a volume group.

- You need at least one unassigned drive.
- Limits exist as to how much drive capacity you can have in a single volume group. These limits vary according to your host type.
- To enable shelf/drawer loss protection, you must create a volume group that uses drives located in at least three shelves or drawers, unless you are using RAID 1, where two shelves/drawers is the minimum.

Review how your choice of RAID level affects the resulting capacity of the volume group.

- If you select RAID 1, you must add two drives at a time to make sure that a mirrored pair is selected. Mirroring and striping (known as RAID 10 or RAID 1+0) is achieved when four or more drives are selected.
- If you select RAID 5, you must add a minimum of three drives to create the volume group.
- If you select RAID 6, you must add a minimum of five drives to create the volume group.

Steps

1. Select **Storage > Pools & Volume Groups**.

2. Click **Create > Volume group**.

The Create Volume Group dialog box appears.

3. Type a name for the volume group.

4. Select the RAID level that best meets your requirements for data storage and protection.

The volume group candidate table appears and displays only the candidates that support the selected RAID level.

5. (Optional) If you have more than one type of drive in your storage array, select the drive type that you want to use.

The volume group candidate table appears and displays only the candidates that support the selected drive type and RAID level.

6. Select the volume group candidate that you want to use based on the following characteristics, and then click **Create**.

Characteristic	Use
Free Capacity	Shows the available capacity in GiB. Select a volume group candidate with the capacity for your application's storage needs.
Total Drives	Shows the number of drives available for this volume group. Select a volume group candidate with the number of drives that you want. The more drives that a volume group contains, the less likely it is that multiple drive failures will cause a critical drive failure in a volume group.

Characteristic	Use
Secure-Capable	<p>Indicates whether this volume group candidate is comprised entirely of secure-capable drives, which can be either Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives.</p> <ul style="list-style-type: none"> You can protect your volume group with Drive Security, but all drives must be secure-capable to use this feature. If you want to create an FDE-only volume group, look for Yes - FDE in the Secure-Capable column. If you want to create a FIPS-only volume group, look for Yes - FIPS in the Secure-Capable column. You can create a volume group comprised of drives that might or might not be secure-capable or are a mix of security levels. If the drives in the volume group include drives that are not secure-capable, you cannot make the volume group secure.
Enable Security?	<p>Provides the option for enabling the Drive Security feature with secure-capable drives. If the volume group is secure-capable and you have set up a security key, you can enable Drive Security by selecting the check box.</p> <p>Note: The only way to remove Drive Security after it is enabled is to delete the volume group and erase the drives.</p>
DA Capable	<p>Indicates if Data Assurance (DA) is available for this group. Data Assurance (DA) checks for and corrects errors that might occur as data is communicated between a host and a storage array.</p> <p>If you want to use DA, select a volume group that is DA capable. This option is available only when the DA feature has been enabled.</p> <p>A volume group can contain drives that are DA-capable or not DA-capable, but all drives must be DA capable for you to use this feature.</p>
Shelf Loss Protection	<p>Shows if shelf loss protection is available.</p> <p>Shelf loss protection guarantees accessibility to the data on the volumes in a volume group if a total loss of communication to a shelf occurs.</p>
Drawer Loss Protection	<p>Shows if drawer loss protection is available, which is provided only if you are using a drive shelf that contains drawers.</p> <p>Drawer loss protection guarantees accessibility to the data on the volumes in a volume group if a total loss of communication occurs with a single drawer in a drive shelf.</p>

Creating a volume group using the AMW

Using SANtricity Storage Manager, you create a volume group, or a logical group of drives. You then designate a portion of the volume group as a volume to present to the host.

About this task

If you are using the Drive Security premium feature, make sure you understand how to implement it. For details, search for the Drive Security topic in the SANtricity Storage Manager Online Help.

Steps

1. Verify that hot spare coverage is adequate for the storage array.
 - a. From the **Array Management Window**, select **Hardware > Hot Spare Coverage**.
 - b. On the **Hot Spare Drive Options** dialog box, select **View/change current hot spare coverage** and select **OK**.
 - c. On the **Hot Spare Coverage** dialog box, view coverage to determine if you need to select more drives for hot spares.

Note: For help determining if coverage is adequate, select the hyperlink “*Tips on providing hot spare coverage*” on the Hot Spare Coverage dialog box.
 - d. If coverage is inadequate, select the **Assign** button and select hot spare drives on the **Assign Hot Spare** dialog box.
 - e. Select **Close**.
2. Select the **Storage & Copy Services** tab, right-click **Total Unconfigured Capacity**, and then select **Create Volume Group**.

Note: If there is more than one drive type, such as SAS and SSD drives, you cannot create a volume group from the high-level **Total Unconfigured Capacity** object. Instead, you must select a sub-object under that high-level object.
3. On the **Introduction** page of the wizard, select **Next**.
4. On the **Volume Group Name & Drive Selection** page of the wizard, perform the following steps:
 - a. Enter a name for the new volume group.
 - b. Select the **Automatic (Recommended)** radio button from the **Drive selection choices** list, and then select **Next**.
5. On the **RAID Level and Capacity** page, perform the following steps:
 - a. Select the desired RAID level for the new volume group from the drop-down list.

Note: For help determining the best RAID level, select the hyperlinks “*What RAID level is best for my application?*” and “*What is tray loss protection?*” on the RAID Level and Capacity page.
 - b. Select the desired volume group configuration from the list of available configurations and select **Finish**.
 - c. The **volume group** wizard automatically displays a prompt for you to create a volume in the newly created volume group. To create a volume immediately, select **Yes** to continue with the volume creation.

Host-side storage considerations

In addition to the express methods for discovering, configuring, and verifying storage on the host, you can explore alternate power methods.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Linux Express Guide](#)

[SANtricity System Manager 11.40 Installing and Configuring for Linux Express Guide](#)

Storage partitions

A storage partition is a logical entity that consists of one or more volumes that can be accessed by a single host or can be shared among hosts that are part of a host group. A host group is a group (cluster) of two or more hosts that share access, in a storage partition, to specific volumes on the storage array. You can create an optional logical entity in the storage management software. You must create a host group only if you will use storage partitions.


Note: If you have an E2800 controller shelf, storage partitioning is neither available nor needed on your system.

Note: If you must define a host group, you can define it through the Define Hosts Wizard described in the AMW online help.


- You can think of a storage partition as a virtual storage array. That is, take the physical storage array and divide it up into multiple virtual storage arrays that you can then restrict to be accessible only by certain hosts.
- You do not create storage partitions in this step, but you must understand them to define your hosts.
- Even if you do not use storage partitions, you must select the Host Operating System type for the Default Group.
- You *do not* need to create storage partitions if these conditions exist:
 - You have only one attached host that accesses all of the volumes on the storage array.
 - You plan to have all of the attached hosts share access to all of the volumes in the storage array.

Note: When you have multiple hosts accessing the volumes in a storage partition, you must have some type of clustering software on the hosts to manage volume sharing and accessibility.

The following displays an example of no additional storage partitions required:




A single host accesses **all** volumes;
no additional storage partitions are needed.



Multiple homogeneous hosts share access to **all** volumes;
no additional storage partitions are needed and
no specific host group is needed.

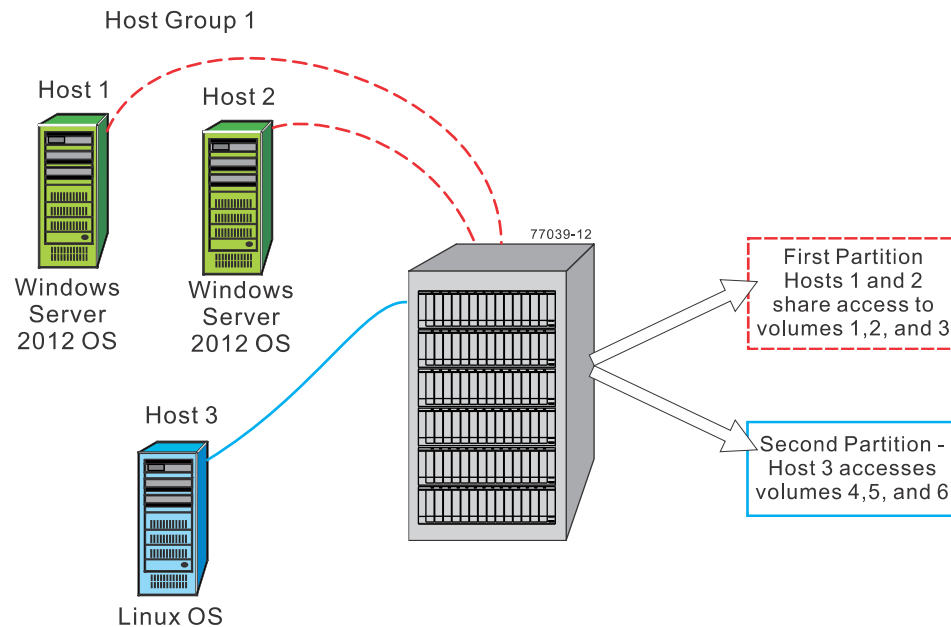
- You *do* need to create storage partitions if these conditions exist:
 - You want certain hosts to access only certain volumes.

The following displays an example of additional storage partitions required (homogeneous host):



- Each host needs access to specific volumes.
 - Both hosts use the same operating system (homogeneous).
 - Storage divided into two logical storage partitions.
 - A Default Group (partition) is not used.
- You have hosts with different operating systems (heterogeneous) attached in the same storage array. You must create a storage partition for each type of host.

The following displays an example of additional storage partitions required (heterogeneous host):



- Host 1 and host 2 (Windows Server 2012 OS) share access to specific volumes through host group 1.
- Two heterogeneous hosts (Linux OS and Windows Server 2012 OS) exist.
- Host 3 (Linux) accesses specific volumes.
- Storage is divided into two logical storage partitions.
- A Default Group (partition) is not used.

Using the `iscsiadm open-iscsi` administration utility with E-Series products

iSCSI connections on Linux use the `open-iscsi` utilities. One command provided by `open-iscsi` encapsulates its capabilities, `iscsiadm`. Because of the complex flags used to describe ifaces, discovered targets, and sessions, the command is complex. You should consult the manual (`man`) page.

The following provides basic information and examples for the common mode (`-m` flag) options:

- `iface` -- configuring ifaces
- `discovery` -- start discovery
- `session` -- login or logout of sessions
- `node` -- display node records

iface

iSCSI `iface` is an alias to a physical NIC port which will be used for iSCSI traffic. Additional TCP parameters can be set in the `iface` command. iSCSI ifaces might be

required in some configurations to guarantee traffic is routed through the appropriate physical interface.

To display all ifaces within a system, specify the following:

```
iscsiadm -m iface
```

Note that the output resulting from this command can be difficult to read.

To display selected ifaces, use the `-I` parameter:

```
iscsiadm -m iface -I myiface01
```

To create a new iface (`myiface01`), use the following command:

```
iscsiadm -m iface -I myiface01 -o new
```

To log out a particular iface (`myiface01`), use the following command:

```
iscsiadm -m iface -I myiface01 -u
```

To delete a particular iface (`myiface01`), use the following command:

```
iscsiadm -m iface -I myiface01 -o delete
```

Required fields for each iface include the hardware address (MAC), the ip address, the transport name (if using a hardware initiator), and the initiator name. All of these options can be configured in `/etc/iscsi/iface` where the file actually exists, or preferably through the `iscsiadm` command. Most information can be taken from the `ifconfig` command or from the `/var/lib/iscsi/` (or on older systems `/etc/iscsi/`) directory.

The following example displays syntax for displaying the hardware address, the ip address, and the initiator name:

```
# iscsiadm -m iface -I myiface01 -o update -n iface.hwaddress -v
00:11:22:33:44:55
# iscsiadm -m iface -I myiface01 -o update -n iface.ipaddress -v
192.168.0.0
# iscsiadm -m iface -I myiface01 -o update -n
iface.initiatorname -v iqn.1994-05.com.redhat:abcdef123456
```

Discovery

Use the discovery mode to send `targets` command output to the target ip. Discovery reports all available target port addresses including IPv4 and IPv6 (if enabled). It also includes target IP's that are not in the same subnet as the target port used for discovery. Discovery is required for automatic login (`-L`), but not for manual login (`-l`). Copy down the target `iqn`, because it is required for login.

```
# iscsiadm -m discovery -t st -p <target ip> -I myiface01
```

Discovery output should resemble `192.168.0.0:3260,2 iqn.1994-08.com.netapp.`

Login

Login can occur to the addresses discovered from the Discovery flag. Manual login is a best practice (-l). Recall the target iqn from the discovery.

Login to the target iqn from myiface01:

```
iscsiadm -m node -T <target iqn> -p <target ip> -I myiface01 -l
```

Use the following command to display current active sessions:

```
iscsiadm -m session
```

Logout

Use the following commands to logout of iSCSI sessions:

Log out an individual iSCSI target by specifying:

```
iscsiadm -m node -u -I <target iqn> -p <target ip>:3260
```

Log out all iSCSI targets by specifying:

```
iscsiadm -m node -u
```

Delete

Use the following command to clear the discovered IPs on a particular target IQN:

```
iscsiadm -m node -u -T <target iqn> -p <target ip>:3260 -o delete
```

Linux partitions, filesystems, and mounting recommendations

The Express guide provides direct commands to partition (`parted`), create a filesystem (`mkfs`), and mount the filesystem for validating access to storage. There are more critical administration options you might want to consider.

The `parted` command can be used to quickly partition a disk, but an administrator should consider using a volume manager on Linux such as LVM. Using a volume manager makes managing complex configurations less complex, but has an upfront learning curve. The filesystem chosen should also be considered. As an example, EXT4 is a very common filesystem on Linux and is used in the Express Guide, but is not the best choice for all deployments. E-Series supports the natively-provided filesystems on each distribution. This information is provided on [Interoperability Matrix Tool](#) in the E-Series Policy and Guidelines.

For persistent deployments, administrators should consider the following flags when adding a mount into `/etc/fstab`:

nobarrier

Allows data to sit in cache instead of being flushed. There is a large performance gain on particular workloads by allowing `nobarrier`. This option should only be used for E-Series storage, because internal disks might not have battery backup.

_netdev

Required for configurations using iSCSI and iSER network protocols. The `_netdev` option forces the mount to wait until the network is up before trying to mount. Without

this option, the os attempts to mount the disk prior to the network being completely available, and it could lead to various timeouts or the os entering recovery mode.

discard

If the storage volume is thinly provisioned, providing the discard flag allows the filesystem to reclaim space. This flag can cause performance degradation. Administrators who want to control when discards take place (for example, nightly), should consider using `fstrim` or an equivalent command for the os.

noatime

Forces file reads to not record their access times to disk, which can increase I/O dramatically on heavy read loads. Setting the `noatime` flag is only recommended for filesystems or dependent applications where a record of the last access time of a file for reading is unnecessary.

Services on Linux that apply to E-Series products

Distributions differ in how services are managed. Some will use the `init` system in which programs are determined by the concept of a run level. Other distributions will use `systemd` to control services. Common commands and equivalencies are provided below.

The following provides common commands and their equivalent across both `init` systems and `systemd` systems.

Command description	init	systemd
view	<code>chkconfig --list</code>	<code>systemctl -t service</code>
start	<code>service <name> start</code>	<code>systemctl start <name></code>
stop	<code>service <name> stop</code>	<code>systemctl stop</code>
restart	<code>service <name> restart</code>	<code>systemctl restart <name></code>
persistently enable	<code>chkconfig <name> on</code>	<code>systemctl enable <name></code>
persistently disable	<code>chkconfig <name> off</code>	<code>systemctl disable <name></code>

Copyright information

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

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