



SANtricity® 11.40

Installing and Configuring for Windows®

Power Guide for Advanced Users

September 2017 | 215-11888_A0
doccomments@netapp.com

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
Windows Server Core: Installing SANtricity Storage Manager	16
Installing the storage array as a boot device	17
Installing SANtricity Storage Manager packages using silent mode	19
Deciding which packages to install	19
Host operating systems	19
Storage management software components	19
Installing the SANtricity software on hosts, monitors, and management stations	21
Adding the storage array to the management domain	27
Preparing to add the storage array to the management domain	27
Completing preliminary tasks for preparing the storage array	27
Setting IP addresses	27
Naming the storage array	28
Passwords	29
Choosing the method for adding the storage array to the management domain	30
Configuring management port IP addresses using the Quick Connect utility	31
Using automatic discovery to add storage arrays to the management domain	31
Manually configuring the controllers by setting up a temporary private network	32
Configuring management port using System Manager	34
Configuring a management port using Storage Manager	35
Configuring multipath	37
Overview of multipath drivers	37
Multipath driver setup considerations	37
Supported multipath drivers	38
Automatic Load Balancing feature overview	39
Multipath configuration diagrams	40
How a multipath driver responds to a data path failure	42
User responses to a data path failure	43
Failover drivers for the Windows operating system	43
Terminology	44
Operational behavior	44

Administrative and configuration interfaces	49
Error handling and event notification	50
Understanding the dsmUtil utility	55
Windows multipath DSM event tracing and event logging	57
Power methods for configuring multipath	61
Dividing I/O activity between two RAID controllers to obtain the best performance	61
Installing the multipath software	62
Compatibility and migration	63
Installation and removal	64
Configuring host utilities, virtualization, and clustering	68
Virtualization considerations	68
Multipathing and virtualization	69
Host clustering support	69
Cluster accessibility	70
Cluster topology	71
Cluster shared storage in SANtricity	71
What are SCSI reservations?	72
Deciding whether to use disk pools or volume groups	73
Creating a volume group	73
Creating a volume group using the AMW	76
Storage partitions	77
Copyright information	81
Trademark information	82
How to send comments about documentation and receive update notifications	83

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	<p>Consider adding options or modifying your installation:</p> <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

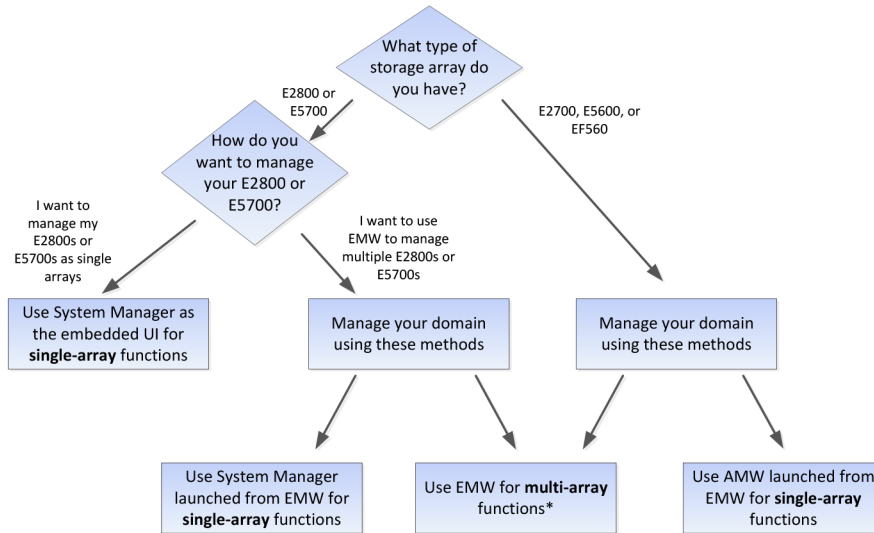
Your options for storage management software vary depending on the array type.

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.

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 21.



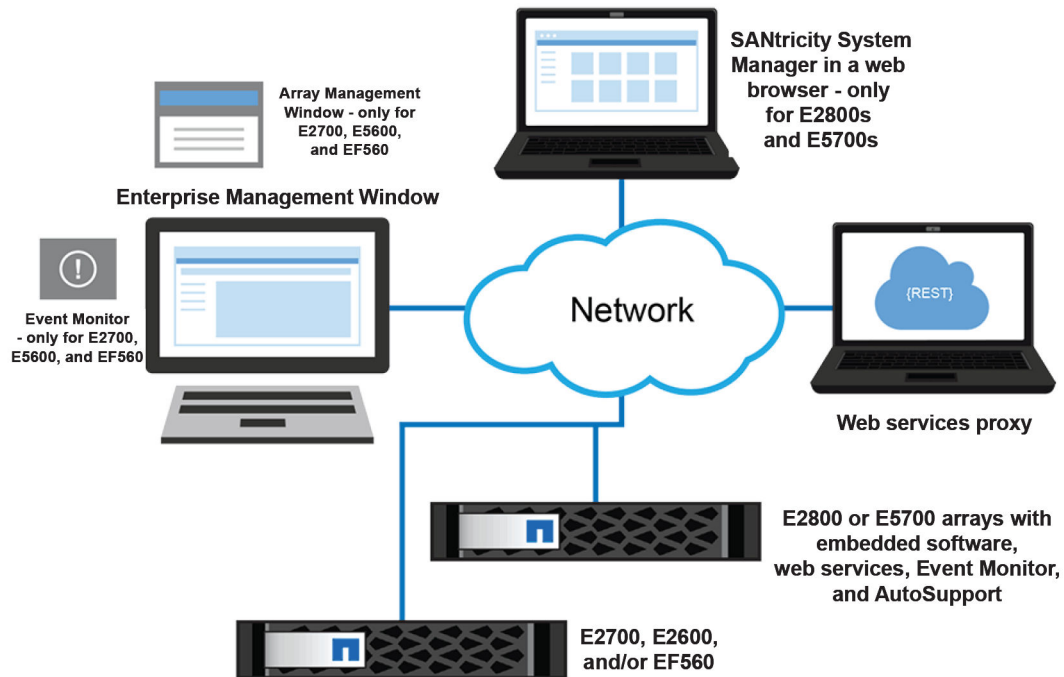
- Multiple E2800 or E5700 storage arrays** — If you have more than one E2800 or E5700 storage array, you can install the EMW to manage your storage environment while handling storage array-based configuration through SANtricity System Manager. The EMW is included with SANtricity Storage Manager.



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 Windows Express Guide](#)
[SANtricity System Manager 11.40 Installing and Configuring for Windows 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 19.</p>
Using the storage array as a boot device	No	Yes See Installing the storage array as a boot device on page 17.

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 19.
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 27 and <i>Choosing the method for adding the storage array to the management domain</i> on page 30.
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 73.

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
Out-of-band <i>with</i> a DHCP server (IPv4 networks only)	<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">  </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>	No additional disadvantages.

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 19.)</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 Windows Express Guide](#)

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

Windows Server Core: Installing SANtricity Storage Manager

Before installing SANtricity Storage Manager, you must first obtain an installation file that is specific to your operating system and to release level of the storage management software.

About this task

In the following steps, the installation file is identified as SMIA-WINX64-11.30.nnnn.nnnn.exe. The characters nnnn.nnnn represent alpha-numeric characters. For example, SMIA-WINx64-11.30.0300.0020.exe. Download this file from [NetApp Support](#).

Steps

1. Download or copy the installation file, SMIA-WINx64-11.30.nnnn.nnnn.exe, to a directory on your host.
2. You have three options for doing the installation:
 - You can specify the `console` parameter during the installation, for example:

```
<hsw executable
.exe> -i console
```

Using this option, questions appear on the console that enable you to choose installation variables. This installation does not use a graphical user interface (GUI). Contact your technical support if you need to change the installation options.

- You can specify the `silent` parameter during the installation, for example:

```
<hsw executable
.exe> -i silent
```

Using this option, the command installs the storage management software using all of the defaults. A silent installation uses a resource file that contains all of the required information, and it does not return any windows until the installation is complete. This installation does not use a graphical user interface (GUI). Contact technical support if you need to change the installation options.

- You can use the **SANtricity InstallAnywhere** installation. For example:

```
<hsw executable
.exe>
```

3. Change your current directory to the installation directory by typing `cd <install directory>` on the command line and then select **Enter**.
4. Type `SMIA-WINx64-11.30.nnnn.nnnn.exe`.
The SANtricity Storage Manager InstallAnywhere dialog is displayed.
5. Press the **Next** button on the first screen of the **SANtricity Storage Manager InstallAnywhere** wizard.
The license agreement is displayed on the second screen of the wizard.
6. Select the **I accept the terms of the license agreement** radio button, then select **Next**.
The Select Installation Type screen is displayed.
7. On the **Select Installation Type** screen, select the **Typical (Full Installation)** radio button. This choice installs both the SMclient software as well as the MPIO multipath driver. Select **Next**.
The Automatically Start Monitor? screen is displayed.
8. On the **Automatically Start Monitor?** screen, select either the **Automatically start monitor** or the **Do not automatically start the monitor** radio button, then select **Next**.
The Pre-Installation Summary screen is displayed. Make note of the Install directory where the software will reside.
9. On the **Pre-Installation Summary** screen, select the **Yes, restart my system** radio button. Then select **Done**.
10. After the system reboots, make sure that the appropriate files are listed in the installation directory (for example: `C:\ProgramFiles (x86)\StorageManager`).
A full installation should include these directories:
 - `util` (SMutil)
 - `client` (SMclient)
 - `agent` (SMagent)
11. Change to the `client` directory and type `SMclient.bat`.

Related concepts

[Configuring multipath](#) on page 37

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:

Windows: `SMIA.xx.xx.xx.xx.exe -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

This package contains the multipath driver that 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 Windows Express Guide](#)
[SANtricity System Manager 11.40 Installing and Configuring for VMware 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):

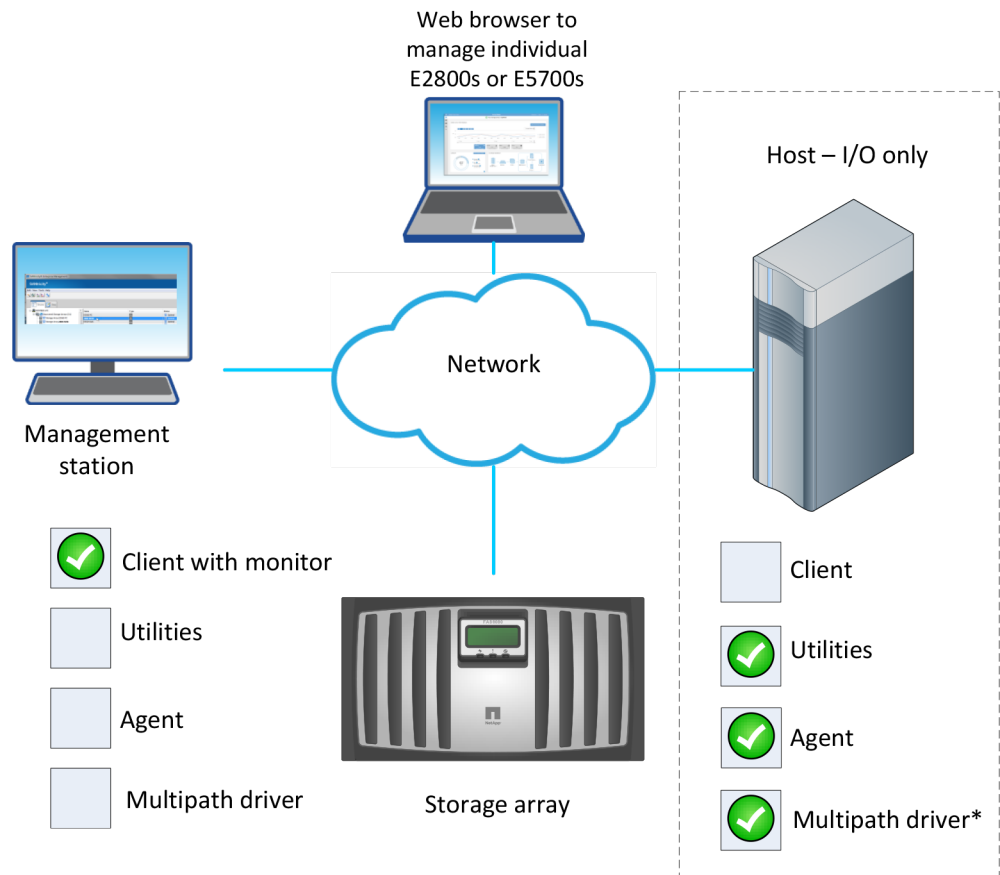
A multipath package is included in the SANtricity installer.

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

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

Installing on the host (I/O only)

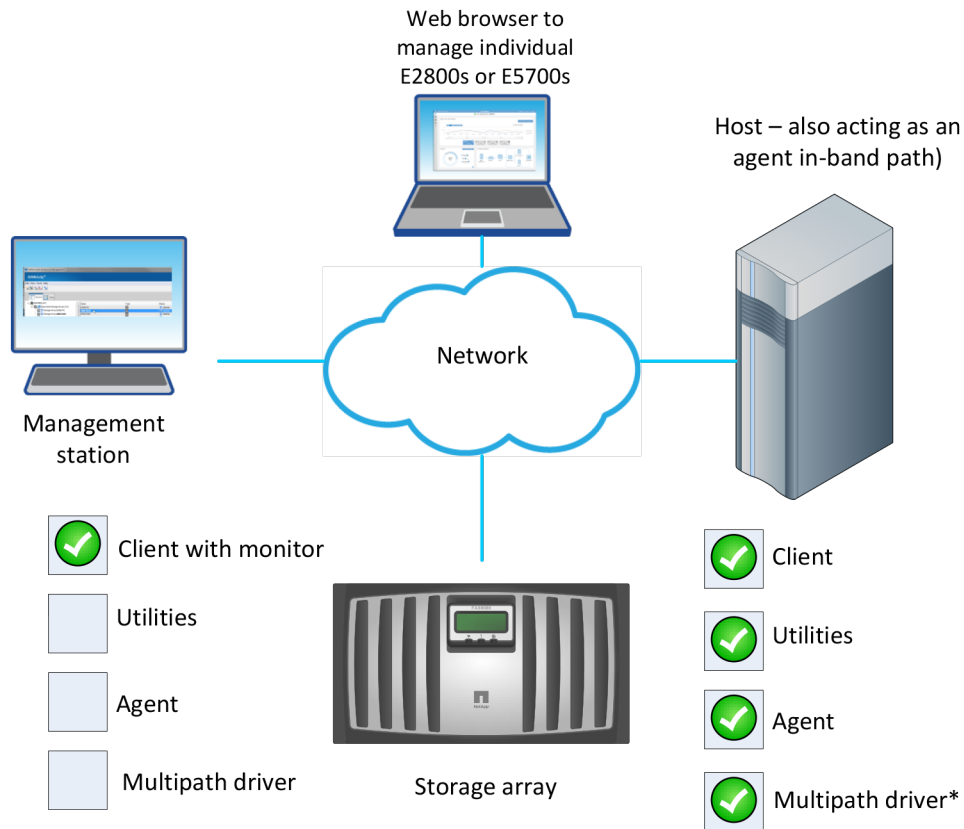
The following conceptual diagram and table provide basic information for installing on the host only for I/O.



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">UtilitiesAgentMultipath driver*	Host	
*The Multipath driver is part of the SANtricity Host installation package for both Windows and AIX operating systems.		

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

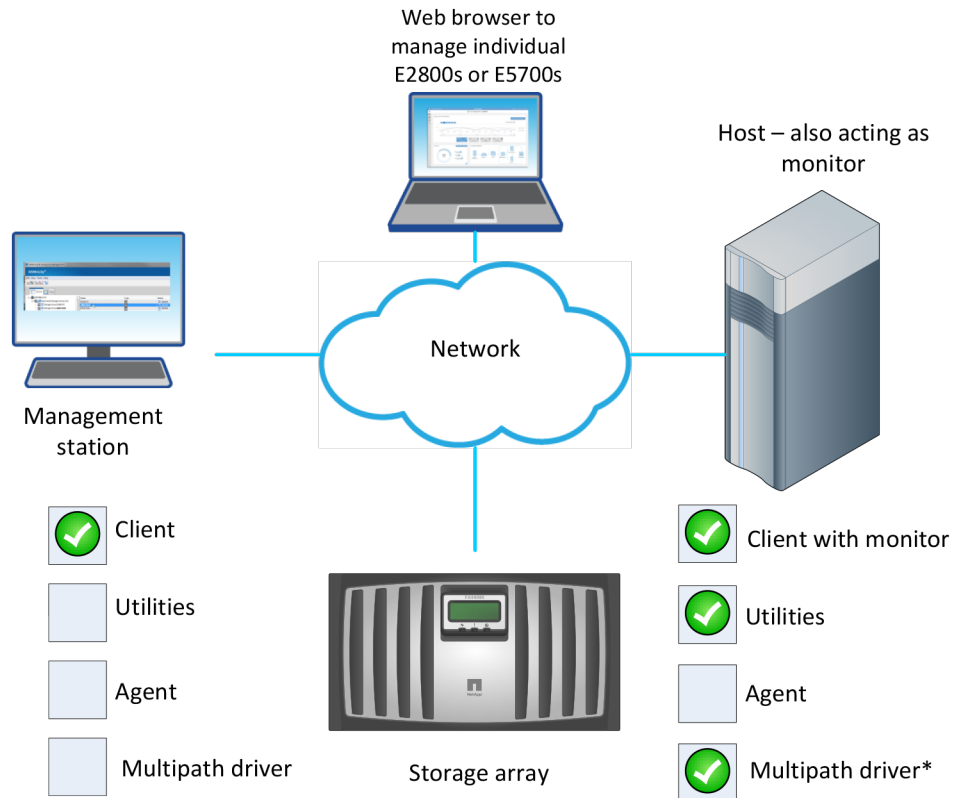
The following conceptual diagram and table provide basic information for installing the host for in-band management.



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?
*If you are managing a single E2800 or E5700 storage array, you do not need a separate Management station.			

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 30

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 31.
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 31. Note: For E2700 and E5600 arrays only, the management station must reside on the same subnetwork as the array during controller management IP configuration.
Stateless address auto-configuration, no DHCP server	
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 32.

Related information

[SANtricity Storage Manager 11.40 Installing and Configuring for Windows Express Guide](#)
[SANtricity System Manager 11.40 Installing and Configuring for Windows 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).
- You have turned on the legacy management interface (SYMbol). If you have disabled the interface, see the *SANtricity System Manager online help* or the *Command Line Reference* for information on re-enabling it.

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 34.

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.

- In the **EMW Setup** tab, select **Name/Rename Storage Arrays**.
- In the **Select storage array** list, select the storage array you added.
- 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 = 0mx, 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 34.
 - For E2700 and E5600 storage arrays, see [Configuring a management port using SANtricity Storage Manager](#) on page 35.
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.

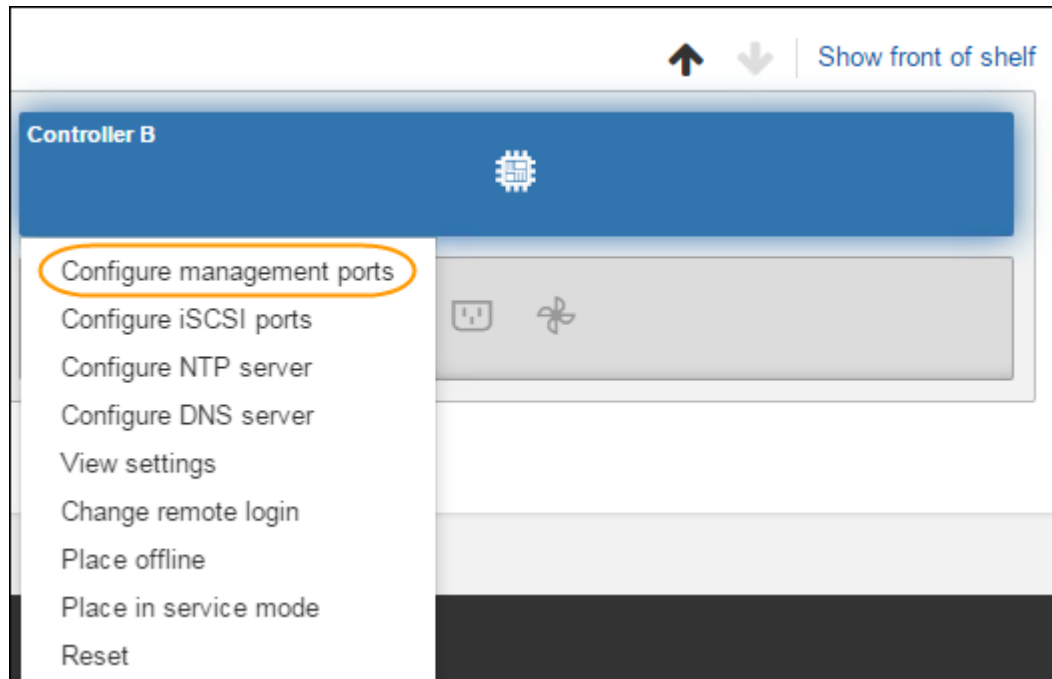
About this task

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.
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

About this task

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 61

[Configuring host utilities, virtualization, and clustering](#) on page 68

Related information

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

[SANtricity System Manager 11.40 Installing and Configuring for Windows 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
Windows Server	MPIO with NetApp E-Series Device Specific Module (DSM) (with ALUA support)	Windows or Windows Clustered
Windows	ATTO Multipath Director	Windows ATTO Note: You must use ATTO FC HBAs.
Windows	ATTO Multipath Director and clustered/parallel file system	ATTO Cluster/All OS Note: You must use ATTO FC HBAs and clustered/parallel file systems.

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 multipath driver is installed as part of the SANtricity Host install package.

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 19.

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.

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
VMware	Native Multipathing Plugin (NMP) with <code>VMW_SATP_ALUA</code> 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, or iSCSI Host Bus Adapters
2. SAS, Fibre Channel, or iSCSI 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.

Failover drivers for the Windows operating system

The failover driver for hosts with Microsoft Windows operating systems is Microsoft Multipath I/O (MPIO) with a Device Specific Module (DSM) for SANtricity Storage Manager.

Enabling and disabling Automatic Load Balancing

Automatic Load Balancing is enabled by default on all storage arrays that ship with SANtricity OS (controller software) 8.30 or later. Linux, Windows, and VMware multipath drivers can use the Automatic Load Balancing feature. If you upgrade your controller from SANtricity OS (controller software) 8.25 or earlier to 8.30 or later, Automatic Load Balancing is disabled by default on the storage array.

You might want to disable Automatic Load Balancing on your storage array for the following reasons:

- You do not want to automatically change a particular volume's controller ownership to balance workload.
- You are operating in a highly tuned environment where load distribution is purposefully set up to achieve a specific distribution between the controllers.

In SANtricity Storage Manager, select the **Storage Array > Configuration > Automatic Load Balancing** menu option to enable or disable the Automatic Load Balancing feature for an individual storage array.

In SANtricity System Manager, select **Settings > System**, scroll down to the Additional Settings section, click the **Enable/Disable Automatic Load Balancing** link, and select the **Enable/Disable automatic load balancing** checkbox to enable or disable the feature for an individual storage array.

Terminology

The Device Specific Module (DSM) for SANtricity Storage Manager uses a generic data model to represent storage instances and uses the following terminology.

- `DeviceInfo` - A specific instance of a logical unit mapped from a storage array to the host that is visible on an I-T nexus.
- `MultipathDevice` - An aggregation of all `DeviceInfo` instances that belong to the same logical unit. Sometimes known as a Pseudo-Lun or Virtual Lun.
- `TargetPort` - A SCSI target device object that represents a connection between the initiator and target (for example, an I-T nexus). This is also known as a Path.
- `TargetPortGroup` - A set of `TargetPort` objects that have the same state and transition from state to state in unison. All `TargetPort` objects associated with a storage array controller belong to the same `TargetPortGroup`, so a `TargetPortGroup` instance can be thought of as representing a Controller.
- `OwningPortGroup` - The `TargetPortGroup` currently being used to process I/O requests for a multi-path device.
- `PreferredPortGroup` - The `TargetPortGroup` that is preferred for processing I/O requests to a multi-path device. The Preferred Port Group and Owning Port Group might be the same or different, depending on the current context. Preferred Port Groups allow for load balancing of multi-path devices across `TargetPortGroups`.
- `PortGroupTransfer` - One or more actions that are necessary to switch the Owning Port Group to another `TargetPortGroup`, for example, to perform failover of one or more LUNs. (Also known as LUN Transfer or Transfer.)

Operational behavior

System environment

Microsoft MPIO is a feature that provides multipath IO support for Windows Operating Systems. It handles OS-specific details necessary for proper discovery and aggregation of all paths exposed by a storage array to a host system. This support relies on built-in or third-party drivers called Device-Specific Modules (DSMs) to handle details of path management such as load balance policies, IO error handling, failover, and management of the DSM.

A disk device is visible to two adapters. Each adapter has its own device stack and presents an instance of the disk device to the port driver (`storport.sys`), which creates a device stack for each instance of the disk. The MS disk driver (`msdisk.sys`) assumes responsibility for claiming ownership of the disk device instances and creates a multipath device. It also determines the correct DSM to use for managing paths to the device. The MPIO driver (`mpio.sys`) manages the connections between the host and the device including power management and PnP management, and acts as a virtual adapter for the multipath devices created by the disk driver.

Failover methods (LUN transfer methods)

The DSM driver supports several different command types ("Methods") of Failover that are described in the next sections.

Mode Select

Mode Select provides a vendor-unique request for an initiator to specify which `TargetPortGroup` should be considered the Owning Port Group.

Target Port Group Support (TPGS)

TPGS provides a standards-based method for monitoring and managing multiple I/O TargetPorts between an initiator and a target. It manages target port states with respect to accessing a DeviceInfo. A given TargetPort can be in different TPGS states for different DeviceInfos. Sets of TargetPorts that have the same state and that transition from state-to-state in unison can be defined as being in the same TargetPortGroup. The following TPGS states are supported.

- **ACTIVE/OPTIMIZED** — TargetPortGroup is available for Read/Write I/O access with optimal performance. This is similar to the concept of a current owning controller.
- **ACTIVE/NON-OPTIMIZED** — TargetPortGroup is available for Read/Write I/O access, but with less than optimal performance.
- **STANDBY** — TargetPortGroup is not available for Read/Write I/O access, but in the event of losing paths to the active TargetPortGroup, this TargetPortGroup can be made available for Read/Write I/O access. This is equivalent to the concept of a non-owning controller.
- **UNAVAILABLE** — TargetPortGroup is not available for Read/Write I/O access and it might not be possible to transition it to a non-UNAVAILABLE state. An example is a hardware failure.

TPGS support is determined by examining the "TPGS" field returned from a SCSI INQUIRY request.

Failover mode

Selective LUN transfers

Selective LUN Transfer is a failover mode that limits the conditions under which the Owning Port Group for a Multipath Device is transferred between TargetPortGroups to one of the following cases:

- Transfer the Multipath Device when the DSM discovers the first TargetPort to the Preferred Port Group.
- Transfer the Multipath Device when the Owning and Preferred Port Group are the same, but the DSM does not have visibility to those groups.
- Transfer the Multipath Device when the DSM has visibility to the Preferred Port Group but not the Owning Port Group.

For the second and third case, configurable parameters exist to define the failover behavior.

Failover method precedence

The Failover method is determined by the DSM on a storage array-by-storage array basis and is based on a system of precedence as described in the following table.

Failover Method	Precedence	Comments
Forced Use of Mode Select	1	Determined by the AlwaysUseLegacyLunFailover configurable parameter. Used when issues are found with TPGS support.
TPGS	2	Determined through a standard Inquiry request.
ModeSelect	3	Default method if all other precedencies are invalidated.

ALUA (I/O shipping)

I/O Shipping is a feature that sends the Host I/O to a Multipath Device to any Port Group within the storage array. If Host I/O is sent to the Owning Port Group, there is no change in existing

functionality. If Host I/O is sent to the Non-Owning Port Group, the SANtricity OS (controller software) uses the back-end storage array channels to send the I/O to Owning Port Group

With I/O Shipping enabled beginning in SANtricity 11.30, most error conditions that require failover results in implicit failback from the controller/target -device. There are, however, cases where failover occurs if the ControllerIoWaitTime is exceeded.

When you install or update the software to SANtricity version 10.83 or later, and install or update the controller SANtricity OS (controller software) to 7.83 or later, support for ALUA is enabled by default.

Path selection (multipath load balancing)

Path selection refers to selecting a TargetPort to a MultipathDevice. When the DSM driver receives a new I/O to process, it begins path selection by trying to find a TargetPort to the Owning Port Group. If a TargetPort to the Owning Port Group cannot be found, and ALUA is not enabled, the DSM driver arranges for MultipathDevice ownership to transfer (or failover) to an alternate TargetPortGroup. The method used to transfer ownership is based on the Failover method defined for the MultipathDevice. When multiple TargetPort's to a MultipathDevice exist, the system uses a load balance policy to determine which TargetPort to use.

Round-Robin with Subset

The Round-Robin with Subset policy selects the most eligible TargetPort in the sequence. TargetPort eligibility is based on a system of precedence, which is a function of DeviceInfo and TargetPortGroup state.

TargetPortGroup State	Precedence
ACTIVE/OPTIMIZED	1
ACTIVE/NON-OPTIMIZED	2
UNAVAILABLE	3
Any other state	Ineligible

Least Queue Depth

The Least Queue Depth policy selects the most eligible TargetPort with the least number of outstanding I/O requests queued. TargetPort eligibility is based on a system of precedence, which is a function of DeviceInfo and TargetPortGroup state. The type of request or number of blocks associated with the request are not considered by the Least Queue Depth policy.

TargetPortGroup State	Precedence
ACTIVE/OPTIMIZED	1
ACTIVE/NON-OPTIMIZED	2
UNAVAILABLE	3
Any other state	Ineligible

Failover Only

The Failover Only policy selects the most eligible TargetPort based on a system of precedence, which is a function of DeviceInfo and TargetPortGroup state. When a TargetPort is selected, it is used for subsequent I/O requests until its state transitions, at which time another TargetPort is selected.

TargetPortGroup State	Precedence
ACTIVE/OPTIMIZED	1

TargetPortGroup State	Precedence
ACTIVE/NON-OPTIMIZED	2
UNAVAILABLE	3
Any other state	Ineligible

Least Path Weight

The Least Path Weight policy selects the most eligible TargetPort based on a system of precedence in which a weight factor is assigned to each TargetPort to a DeviceInfo. I/O requests are routed to the lowest weight TargetPort of the Owning Port Group. If the weight factor is the same between TargetPorts, the Round-Robin load balance policy is used to route I/O requests.

TargetPortGroup State	Precedence
ACTIVE/OPTIMIZED	1
ACTIVE/NON-OPTIMIZED	2
UNAVAILABLE	3
Any other state	Ineligible

Additional Notes On Path Selection

If the only eligible TargetPortGroup states are STANDBY, a Failover Method is initiated to bring the TargetPortGroup state to ACTIVE/OPTIMIZED or ACTIVE/NON-OPTIMIZED.

Online/Offline path states

The ACTIVE/OPTIMIZED and ACTIVE/NON-OPTIMIZED states reported by TargetPortGroup and DeviceInfo objects are from the perspective of the target (storage array). These states do not take into account the overall condition of the TargetPort connections that exist between the initiator and target. For example, a faulty cable or connection might cause many retransmissions of packets at a protocol level, or the target itself might be experiencing high levels of I/O stress. Conditions like these can cause delays in processing or completing I/O requests sent by applications, and does not cause OS-level enumeration activities (- PnP) to be triggered.

The DSM supports the ability to place the DeviceInfo objects that are associated with a TargetPort into an OFFLINE state. An OFFLINE state prevents any I/O requests from being routed to a TargetPort regardless of the actual state of the connection. The OFFLINE state can be performed automatically based on feature-specific criteria (such as Path Congestion Detection). It also can be performed through the multipath utility (dsmUtil) but known as ADMIN_OFFLINE instead. A TargetPort in an ADMIN_OFFLINE state can be placed only in an ONLINE state by an Admin action, host reboot, or PnP removal/add.

Path Congestion Detection

Path Congestion Detection monitors the I/O latency of requests to each TargetPort, and is based on a set of criteria that automatically place the TargetPort into an OFFLINE state. The criteria are defined through configurable parameters.

Example Configuration Settings for the Path Congestion Detection Feature

Note: Before you can enable path congestion detection, you must set the CongestionResponseTime, CongestionTimeFrame, and CongestionSamplingInterval parameters to valid values.

To set the path congestion I/O response time to 10 seconds, do the following:

```
dsmUtil -o CongestionResponseTime=10,SaveSettings
```

To set the path congestion sampling interval to one minute, do the following:

```
dsmUtil -o CongestionSamplingInterval=60,SaveSettings
```

To enable Path Congestion Detection, do the following:

```
dsmUtil -o CongestionDetectionEnabled=0x1,SaveSettings
```

To set a path to Admin Offline, do the following:

```
dsmUtil -o SetPathOffline=0x77070001
```

Note: You can find the path ID (in this example 0x77070001) using the `dsmUtil -g` command.

To set a path Online, do the following:

```
dsmUtil -o SetPathOnline=0x77070001
```

Per-Protocol I/O timeouts

The MS Disk driver must assign an initial I/O timeout value for every non-pass-through request. By default, the timeout value is 10 seconds, although you can override it using the Registry setting called `TimeoutValue`. The timeout value is considered global to all storage that the MS Disk driver manages.

The DSM can adjust the I/O timeout value of Read/Write requests (those requests passed by MPIO into the `DsmLBGetPath()` routine) based on the protocol of the `TargetPort` chosen for the I/O request.

The timeout value for a protocol is defined through configurable parameters.

Wait times

A Wait Time is an elapsed time period that, when expired or exceeded, causes one or more actions to take place. There is no requirement that a resource, such as a kernel timer, manage the time period which would immediately cause execution of the action(s). For example, an I/O Wait Time will establish a start time when the I/O request is first delivered to the DSM driver. The end time establishes when the I/O request is returned. If the time period is exceeded, an action such as Failover, is initiated between `TargetPortGroups`.

All Wait Times defined by the DSM driver are configurable and contain the term "WaitTime" as part of the configuration name. The "Configurable parameters" topic provides a complete list of Wait Times.

SCSI reservations

Windows Server Failover Cluster (WSFC) uses SCSI-3 Reservations, otherwise known as Persistent Reservations (PR), to maintain resource ownership on a node. The DSM is required to perform some special processing of PR's because WSFC is not multipath-aware.

Native SCSI-3 persistent reservations

Windows Server 2008 introduced a change to the reservation mechanism used by the Clustering solution. Instead of using SCSI-2 reservations, Clustering uses SCSI-3 Persistent Reservations, which removes the need for the DSM to handle translations. Even so, some special handling is required for certain PR requests because Cluster itself has no knowledge of the underlying `TargetPorts` for a `MultipathDevice`.

Special circumstances for array brownout conditions

Depending on how long a brownout condition lasts, Persistent Registration information for volumes might be lost. By design, WSFC periodically polls the cluster storage to determine the overall health and availability of the resources. One action performed during this polling is a `PRIN READ KEYS`

request, which returns registration information. Because a brownout can cause blank information to be returned, WSFC interprets this as a loss of access to the disk resource and attempts recovery by first failing the resource and then performing a new arbitration. The arbitration recovery process happens almost immediately after the resource is failed. This situation, along with the PnP timing issue, can result in a failed recovery attempt. You can modify the timing of the recovery process by using the `cluster.exe` command-line tool.

Another option takes advantage of the Active Persist Through Power Loss (APTPL) feature found in Persistent Reservations, which ensures that the registration information persists through brownout or other conditions related to a power failure. APTPL is enabled when a PR REGISTRATION is initially made to the disk resource. You must set this option before PR registration occurs. If you set this option after a PR registration occurs, take the disk resource offline and then bring it back online.

WSFC does not use the APTPL feature but a configurable option is provided in the DSM to enable this feature when a registration is made through the multipath utility.

Note:

The SCSI specification does not provide a means for the initiator to query the target to determine the current APTPL setting. Therefore, any output generated by the multipath utility might not reflect the actual setting.

Implicit controller failback

Implicit failback from the controller rather than through the multipath driver is the new default starting with SANtricity 8.30.

Starting with E-Series SANtricity 8.30, the array SANtricity initiates a failback operations from the controller. Auto Failback through the Preferred TargetPortGroup is no longer supported. This change allows the array SANtricity OS software greater flexibility in balancing I/O load across the controllers, because the host multipath solution responds more readily to implicit ownership changes.

MPIO pass-through

One of MPIO's main responsibilities is to aggregate all DeviceInfo objects into a MultipathDevice, based partially on input from the DSM. By default, the TargetPort chosen for an I/O request is based on current Load Balance Policy. If an application wants to override this behavior and send the request to a specific TargetPort, it must do so using an MPIO pass-through command (`MPIO_PASS_THROUGH_PATH`). This is a special IOCTL with information about which TargetPort to use. A TargetPort can be chosen through one of two of the following methods:

- **PathId** — A Path Identifier, returned to MPIO by the DSM when `DsmSetPath()` is called during PnP Device Discovery.
- **SCSI Address** — A `SCSI_ADDRESS` structure, supplied with the appropriate Bus, Target, and ID information.

Administrative and configuration interfaces

This section describes the Windows Management Instrumentation (WMI) and CLI interfaces.

Windows management instrumentation (WMI)

Windows Management Instrumentation (WMI) is used to manage and monitor Device-Specific Modules (DSMs).

During initialization, the DSM passes WMI entry points and MOF class GUID information to MPIO, which publishes the information to WMI. When MPIO receives a WMI request, it evaluates the embedded GUID information to determine whether to forward the request to the DSM or to keep it with MPIO.

For DSM-defined classes, the appropriate entry point is invoked. MPIO also publishes several MOF classes that the DSM is expected to handle. MOF classes also can have Methods associated with them that can be used to perform the appropriate processing task.

CLI interface -- multipath utility (dsmUtil)

The dsmUtil utility is used with the DSM driver to perform various functions provided by the driver.

Configurable parameters

The DSM driver contains field-configurable parameters that affect its configuration and behavior. You can set these parameters using the multipath utility (dsmUtil). Some of these parameters also can be set through interfaces provided by Microsoft.

Persistence of configurable parameters

Each configuration parameter defined by the DSM has a default value that is hard-coded into the driver source. This default value allows for cases where a particular parameter might have no meaning for a particular customer configuration, or a parameter that needs to assume a default behavior for legacy support purposes, without the need to explicitly define it in non-volatile storage (registry). If a parameter is defined in the registry, the DSM uses that value rather than the hard-coded default.

There might be cases where you want to modify a configurable parameter, but only temporarily. If the host is subsequently rebooted, the value in non-volatile storage is used. By default, any configurable parameter changed by the multipath utility only affects the in-memory representation. The multipath utility can optionally save the changed value to non-volatile storage through an additional command-line argument.

Scope of configurable parameters

A localized configurable parameter is one that can be applied at a scope other than global. Currently the only localized parameter is for load balance policy.

Configurable parameters

Using dsmUtil, you can tune failover performance for your workload. Do this with the assistance of Technical Support. Some of the configurable parameter settings for your configuration are listed in [NetApp Interoperability Matrix Tool](#).

Error handling and event notification

Event logging

Event channels

An Event Channel is a receiver ("sink") that collects events. Some examples of event channels are the Application and System Event Logs. Information in Event Channels can be viewed through several means such as the Windows Event Viewer and `wevtutil.exe` command. The DSM uses a set of custom-defined channels for logging information, found under the "Applications and Services Logs" section of the Windows Event Viewer.

Custom event view

The DSM is delivered with a custom Event Viewer filter that can combine the information from the custom-defined channels with events from the System Event Log. To use the filter, import the view from the Windows Event Viewer.

Event messages

For the DSM, each log message is well-defined and contains one or more required `ComponentNames` as defined. By having a clear definition of the event log output, utilities or other applications and services can query the event logs and parse it for detailed DSM information or use it for troubleshooting purposes. The following tables list the DSM event log messages and also includes the core MPIO messages.

All MPIO-related events are logged to the System Event Log. All DSM-related events are logged to the DSM's custom Operational Event Channel.

Event Message	Event Id (Decimal)	Event Severity
Memory Allocation Error. Memory description information is in the DumpData.	1000	Informational
Queue Request Error. Additional information is in the DumpData.	1001	Informational

Event Message	Event Id (Decimal)	Event Severity
<msg>. Device information is in the DumpData.	1050	Informational
<msg>. TargetPort information is in the DumpData.	1051	Informational
<msg>. TargetPortGroup information is in the DumpData.	1052	Informational
<msg>. MultipathDevice is in the DumpData.	1053	Informational
<msg>. Array information is in the DumpData.	1054	Informational
<msg>.	1055	Informational
<msg>. Device information is in the DumpData.	1056	Warning
<msg>. TargetPort information is in the DumpData.	1057	Warning
<msg>. TargetPortGroup information is in the DumpData.	1058	Warning
<msg>. MultipathDevice information is in the DumpData.	1059	Warning
<msg>. Array information is in the DumpData.	1060	Warning
<msg>.	1061	Warning
<msg>. Device information is in the DumpData.	1062	Error
<msg>. TargetPort information is in the DumpData.	1063	Error
<msg>. TargetPortGroup information is in the DumpData.	1064	Error
<msg>. MultipathDevice information is in the DumpData.	1065	Error
<msg>. Array information is in the DumpData.	1066	Error
<msg>.	1067	Error
<msg>.	1068	Error

Event Message	Event Id (Decimal)	Event Severity
IO Error. More information is in the DumpData.	1100	Informational
IO Request Time Exceeded. More information is in the DumpData.	1101	Informational
IO Throttle Requested to <MPIODisk_n>. More information is in the DumpData.	1102	Informational
IO Resume Requested to <MPIODisk_n>. More information is in the DumpData.	1103	Informational
No Path Available for IO to \Device\MPIODisk	1104	Error
<msg>. More information in the DumpData	1105	Warning
<msg>. More information in the DumpData	1106	Informational
<msg>. More information in the DumpData	1107	Informational

Event Message	Event Id (Decimal)	Event Severity
Failover Request Issued to <MPIODisk_n>. More information is in the DumpData.	1200	Informational
Failover Request Issued Failed to <MPIODisk_n>. More information is in the DumpData.	1201	Error
Failover Request Succeeded to <MPIODisk_n>. More information is in the DumpData.	1202	Informational
Failover Request Failed to <MPIODisk_n>. More information is in the DumpData.	1203	Error
Failover Request Retried to <MPIODisk_n>. More information is in the DumpData.	1204	Informational
Failover Error to <MPIODisk_n>. More information is in the DumpData.	1205	Error
<MPIODisk_n> rebalanced to Preferred Target Port Group (Controller). More information is in the DumpData.	1206	Informational
Rebalance Request Failed to <MPIODisk_n>. More information is in the DumpData.	1207	Error
<MPIODisk_n> transferred due to Load Balance Policy Change. More information is in the DumpData.	1208	Informational
Transfer Due to Load Balance Policy Change Failed for <MPIODisk_n>. More information is in the DumpData.	1209	Error
Rebalance Request issued to <MPIODisk_n>. More information is in the DumpData.	1210	Informational
Rebalance Request Issued Failed to <MPIODisk_n>. Array information is in the DumpData.	1211	Error

Event Message	Event Id (Decimal)	Event Severity
Rebalance Request Retried to <MPIODisk_n>. More information is in the DumpData.	1212	Informational
Failover Request Issued to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1213	Informational
Failover Request Issued Failed to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1214	Error
Failover Request Failed to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1215	Error
Failover Request Retried to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1216	Informational
Failover Setup Error for Failover to TargetPortGroup (Controller <n>). More information is in the DumpData.	1217	Error
Failover Request Succeeded to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1218	Informational
Rebalance Request issued to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1219	Informational
Rebalance Request Issued Failed to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1220	Error
Rebalance Request Retried to TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1221	Informational
Rebalance Setup Error for Rebalance to TargetPortGroup (Controller <n>). More information is in the DumpData.	1222	Error
<MPIODisk_n> transferred from TargetPortGroup (Controller <n>) due to Load Balance Policy Change. More information is in the DumpData.	1223	Informational
Transfer Due to Load Balance Policy Change Failed for TargetPortGroup (Controller <n>) via <MPIODisk_n>. More information is in the DumpData.	1224	Error
<MPIODisk_n> rebalance to Preferred TargetPortGroup (Controller <n>). More information is in the DumpData.	1225	Informational
Failure during transfer to TargetPortGroup (Controller <n>). More information is in the DumpData.	1226	Error

Event Message	Event Id (Decimal)	Event Severity
Transfer Setup Due to Load Balance Policy Change Failed for TargetPortGroup (Controller <n>). More information is in the DumpData.	1227	Error

Event Message	Event Id (Decimal)	Event Severity
Configured Parameter Invalid of Out of Range. More information is in the DumpData.	1300	Informational
Configuration Initialization Error	1301	Informational
No Target Ports Found for <MPIODisk_n>. More information is in the DumpData.	1302	Error

Architecture Note:

Event Message	Event Id (Decimal)	Severity
New Device Detected. More information is in the DumpData.	1450	Informational
Device for <MPIODisk_n> Pending Removed via MPIO. More information is in the DumpData.	1451	Informational
Device for <MPIODisk_n> Removed via MPIO. More information is in the DumpData.	1452	Informational
Early Device Failure. More information is in the DumpData.	1453	Warning
Failed to obtain the \Device\MPIODisk ordinal. More information is in the DumpData.	1454	Warning

Event Message	Event Id (Decimal)	Severity
New TargetPort (Path) Detected. More information is in the DumpData.	1600	Informational
TargetPort (Path) Removed via MPIO. More information is in the DumpData.	1601	Informational
TargetPort (Path) Offline Manually. More information is in the DumpData.	1602	Warning
TargetPort (Path) Online Manually. More information is found in the DumpData.	1603	Warning
TargetPort (Path) Offline (Threshold Exceeded). More information is found in the DumpData.	1604	Warning
Congestion Threshold Detected on TargetPort. More information is found in the DumpData.	1605	Warning
Not all PCD configuration parameters are set. PCD is not enabled.	1606	Warning

Event Message	Event Id (Decimal)	Severity
Congestion Threshold detected but path not placed Offline due to configuration setting. More information is found in the DumpData.	1607	Warning
TargetPort (Path) automatically placed Offline due to exceeding congestion threshold. More information is in the DumpData.	1608	Warning

Event Message	Event Id (Decimal)	Severity
New TargetPortGroup (Controller) Detected. More information is in the DumpData.	1750	Informational
TargetPortGroup (Controller) Removed. More information is in the DumpData.	1751	Informational
TargetPortGroup (Controller) IO Timeout. More information is in the DumpData	1752	Error

Event Message	Event Id (Decimal)	Severity
New Storage Array Detected. More information is in the DumpData.	1900	Informational
Storage Array Removed. More information is in the DumpData.	1901	Informational

Understanding the dsmUtil utility

The DSM solution bundles a command-line multipath utility, named dsmUtil, to handle various management and configuration tasks. Each task is controlled through arguments on the command-line.

Reporting

The dsmUtil utility offers the following reporting options.

- **Storage Array Summary ('-a' option)** - Provides a summary of all storage arrays recognized by the DSM, and is available through the `-a` command-line option. For example, to retrieve a summary of all recognized storage arrays use the following command:

```
C:\> dsmUtil -a
```

- **Storage Array Detail ('-a' or '-g' option)** - Provides a detailed summary of multipath devices and target ports for an array, and is available through the `-g` command-line option. The same detailed summary information is also available with an optional argument to `-a`. In either case, the array WWN is specified to obtain the detailed information as shown in the following example:

```
C:\> dsmUtil -a 600a0b8000254d370000000046aaaa4c
```

- **Storage Array Detail Extended ('-a' or '-g' option)** - Extended information, providing further details of the configuration, is available by appending the keyword `extended` to the command-line for either `-a` or `-g` options. Extended information is typically used to assist in troubleshooting

issues with a configuration. Extended information appears as italic but is printed as normal text output.

- **Storage Array Real-Time Status ('-S' option)** - A real-time status of the target ports between a host and array is available using the `-S` command-line option.
- **Cleanup of Status Information ('-c' option)** - Information obtained while running the `-S` option is persisted across host and array reboots. This might result in subsequent calls to the `-S` option producing erroneous results if the configuration has permanently changed. For example, a storage array is permanently removed because it is no longer needed. You can clear the persistent information using the `-c` command-line option.
- **MPIO Disk to Physical Drive Mappings ('-M' option)** - This report allows a user to cross-reference the MPIO Virtual Disk and Physical Disk instance with information from the storage array on the mapped volume. The output is similar to the `smdevices` utility from the SANtricity package.

Administrative and Configuration Interfaces

The `dsmUtil` utility offers the following administrative and configuration interface options.

- **Setting of DSM Feature Options** - Feature Options is an interface exposed by the DSM, through WMI, which can be used for several configuration parameter-related tasks. The `'-o'` command-line option is used to carry out these tasks. Several sub-options are available when using the `'-o'` option for parameter-specific purposes:
 - **Parameter Listing** - If the user specifies no arguments to `'-o'` the DSM returns a list of parameters that can be changed.
 - **Change a Parameter** - If the user requests a parameter value change, the DSM verifies the new parameter value, and if within range applies the value to the parameter. If the value is out of range, the DSM returns an out-of-range error condition, and `dsmUtil` shows an appropriate error message to the user. Note this parameter value change is in-memory only. That is, the change does not persist across a host reboot. If the user wants the change to persist, the `SaveSettings` option must be provided on the command-line, after all parameters have been specified.
- **Setting of MPIO-Specific Parameter** - As originally written, MPIO provided several configuration settings which were considered global to all DSMs. An enhancement was later introduced which applied some of these settings on a per-DSM basis. These settings (global and per-DSM) can be manually changed in the Registry but does not take effect until the next host reboot. They also can take effect immediately, but require that a WMI method from a DSM-provided class is executed. For per-DSM settings, MPIO looks in the `\\HKLM\System\CurrentControlSet\Services\<DSMName>\Parameters` subkey. The DSM cannot invoke MPIO's WMI method to apply new per-DSM settings, therefore `dsmUtil` must do this. The `'-P'` option is used for several tasks related to MPIO's per-DSM setting.
 - **Parameter Listing** - An optional argument to `'-P'` (`GetMpioParameters`) is specified to retrieve the MPIO specific per-DSM settings. All of the MPIO specific settings are displayed to the user as one line in the command output.
 - **Change a Parameter** - If the user requests a parameter value change they provide the parameter name and new value in a `'key=value'` format. Multiple parameters might be issued with a comma between each key/value statement. It appears MPIO does not do any validation of the data passed in, and the change takes effect immediately and persist across reboots.
- **Removing Device-Specific Settings** - The `'-R'` option is used to remove any device-specific settings for inactive devices from the Registry. Currently, the only device-specific settings that persist in the Registry are Load Balance Policy.

- **Invocation of Feature Option Actions/Methods** - Feature Options is an interface exposed by the DSM, through WMI, that also can be used to run specific actions (or methods) within the DSM. An example of an action is setting the state of a TargetPort (ie - path) to Offline. The '-o' command-line option mentioned in the Setting of Feature Options section is used to carry out these tasks. Several sub-options are available when using the '-o' option to run specific actions:
 - Action Listing - If the user specifies no arguments to '-o' the DSM returns a list of actions that can be invoked.
 - Executing An Action - Executing an action is similar to specifying a value for a configuration parameter. The user enters the name of the action, followed by a single argument to the function. The DSM runs the method and returns a success/failure status back to the utility.
- **Requesting Scan Options** - The utility can initiate several scan-related tasks. It uses the '-s' option with an optional argument that specifies the type of scan-related task to perform. Some of these are handled by the DSM while others are handled by the utility.
- **Bus Rescan** - This option causes a PnP re-enumeration to occur, and is invoked using the 'busscan' optional argument. It uses the Win32 configuration management APIs to initiate the rescan process. Communication with the DSM is not required.

Windows multipath DSM event tracing and event logging

The DSM for Windows MPIO uses several methods that you can use to collect information for debugging and troubleshooting purposes. These methods are detailed in this section.

Event tracing

The DSM for Windows MPIO uses several methods to collect information for debugging and troubleshooting purposes. These methods are detailed in this section.

About event tracing

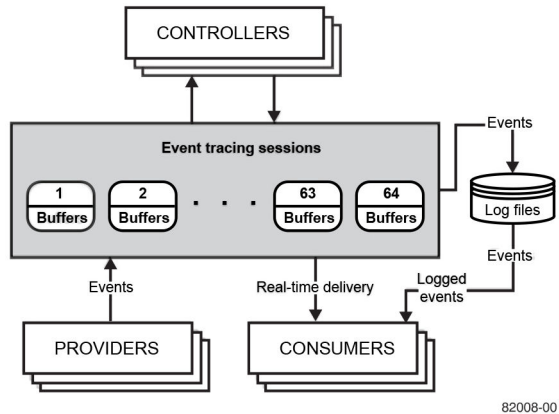
Event Tracing for Windows (ETW) is an efficient kernel-level tracing facility that lets you log kernel or application-defined events to a log file. You can view the events in real time or from a log file and use the events to debug an application or to determine where performance issues are occurring in the application.

ETW lets you enable or disable event tracing dynamically, allowing you to perform detailed tracing in a production environment without requiring computer or application restarts.

The Event Tracing API is divided into three distinct components:

- Controllers, which start and stop an event tracing session and enable providers.
- Providers, which provide the events. The DSM is an example of a Provider.
- Consumers, which consume the events.

The following figure shows the event tracing model.



Controllers

Controllers are applications that define the size and location of the log file, start and stop event tracing sessions, enable providers so they can log events to the session, manage the size of the buffer pool, and obtain execution statistics for sessions. Session statistics include the number of buffers used, the number of buffers delivered, and the number of events and buffers lost.

Providers

Providers are applications that contain event tracing instrumentation. After a provider registers itself, a controller can then enable or disable event tracing in the provider. The provider defines its interpretation of being enabled or disabled. Generally, an enabled provider generates events, while a disabled provider does not. This lets you add event tracing to your application without requiring that it generate events all the time. Although the ETW model separates the controller and provider into separate applications, an application can include both components.

There are two types of providers: the classic provider and the manifest-based provider. The DSM is a classic provider and the tracing events it generates are from the 'TracePrint' API.

Consumers

Consumers are applications that select one or more event tracing sessions as a source of events. A consumer can request events from multiple event tracing sessions simultaneously; the system delivers the events in chronological order. Consumers can receive events stored in log files, or from sessions that deliver events in real time. When processing events, a consumer can specify start and end times, and only events that occur in the specified time frame will be delivered.

What you need to know about event tracing

- Event Tracing uses Non-Paged Pool kernel memory to hold the unflushed events. When configuring trace buffer sizes, try to minimize the buffers potentially used.
- If large trace buffer sizes have been requested at boot, you might experience a delay in boot-time as referenced in this knowledge base article: <http://support.microsoft.com/kb/2251488>.
- If events are being added to the trace buffer faster than can be flushed then you can experience missed events. The logman utility indicates how many events are missed. If you experience this behavior, either increase your trace buffer size or (if flushing to a device) find a device that can handle faster flush rates.

Viewing trace events

Trace events captured to a log file are in a binary format that is not human-readable, but can be decoded properly by technical support. Submit any captured logs to technical support.

Event logging

Windows Event Logging provides applications and the operating system a way to record important software and hardware events. The event logging service can record events from various sources and store them in a single collection called an Event Log. The Event Viewer, found in Windows, enables users to view these logs. Version 1.x of the DSM recorded events in the legacy system log.

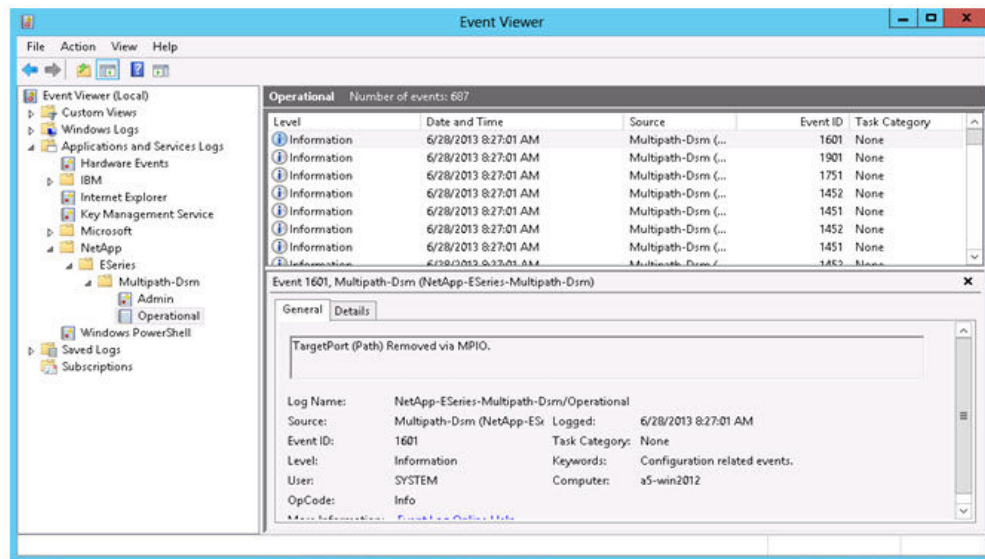
Windows Server 2008 introduced a redesign of the event logging structure that unified the Event Tracing for Windows (ETW) and Event Log APIs. It provides a more robust and powerful mechanism for logging events. Version 2.x of the DSM uses this new approach.

As with Event Tracing, the DSM is considered a provider of Event Log events. Event Log events can be written to the legacy system log, or to new event channels. These event channels are similar in concept to the legacy system log but allow the DSM to record more detailed information about each event generated. In addition, it allows the DSM to record the information into a dedicated log where it won't overwrite or obscure events from other components in the system. Event channels also can support the ability to write events at a higher throughput rate.

Event channels

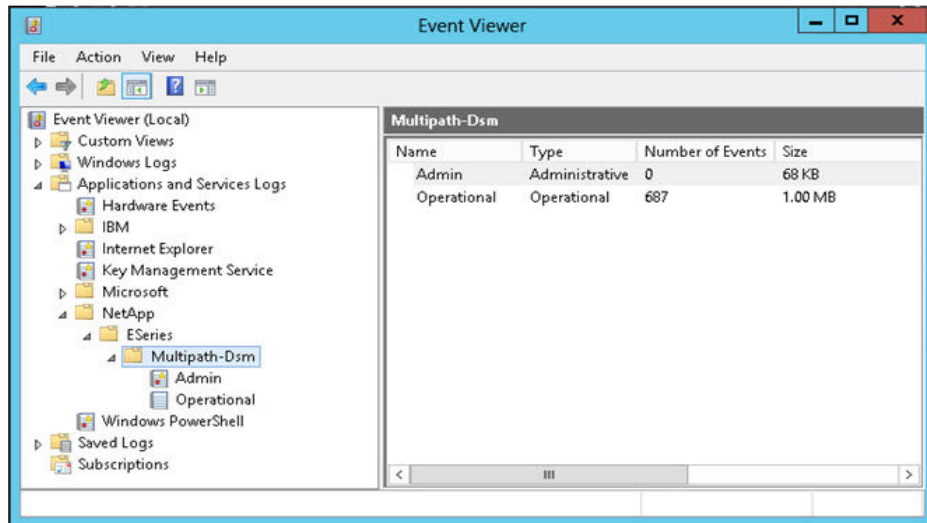
Event channels are viewed using the same Event Viewer application that you use to view the legacy system logs. Currently, the only channel used is the Operational channel.

Events logged into the Admin and Operational channels are stored in the same .EVTX format used by other Windows logs. The following figure shows an example of the event channels.



82008-02

When you select the Operational channel, a tri-pane window appears that shows several rows of events and details of the currently selected event as shown in the following figure. You can select the Details tab to view the raw XML data that makes up the event.



82008-03

Loading the custom event view

The following describes a simple procedure for combining both the DSM and the system log information into a convenient single view.

About this task

You can use the custom view to combine the DSM and system log information into a single view.

Steps

1. In the Event Viewer application, right-click **Custom Views** > **Import Custom View**.
2. Go to the directory where the DSM installation is installed and look in the 'drivers' directory for a file named `CombinedDsmEventChannelView.xml`.
3. Click **OK** to accept the location of the custom view.

A new Custom View named `CombinedDsmEventChannelView` will appear as an option. Select the new custom view to show output from both logs.

Event decoding

Event decoding provides a description of how DSM provides way to store information about an object, and general rules for decoding such information.

Version 2.x of the DSM provides an internally-consistent way of storing information about an object, such as a disk device or controller, which can be provided as part of each record written to an event channel. The component information is a raw stream of bytes that is decoded and merged with the other data to present a complete description of each event record.

1. When the DSM solution is built, the source code is scanned by a script which generates several XML definition files describing details of each Event and the associated base components. These XML definition files are shipped with the solution.
2. Events that need to be decoded are saved to an `.EVTX` file, or can be decoded directly on a Host if there is access to the required Event channels.
3. A PowerShell script and `cmdlet` uses the XML and Event Logs to generate a CSV-formatted document containing the decoded events. This document can be imported to applications such as Excel for viewing.

Files used in the decode process

The 'decoder' directory contains all the files used to decode the event logs.

- `'DecodeEvents.bat'` - This batch file invokes a new powershell session to execute the decoding process. The decoding process will utilize the XML files described below.
- `BaseComponents.xml` - This XML file provides details on each base component and should not be modified as any change can cause a failure in properly decoding events.
- `EventComponents.xml` - This XML file provides details for each event generated by the DSM and the base component data reported. It should not be modified as any change can cause a failure in properly decoding events.
- `LogsToDecode.xml` - This XML file defines the source(s) of the event log data. For convenience the decoding process will not only attempt to decode messages from the DSM, but also messages reported by Microsoft MPIO. This file can be modified as needed to define the location of event log data to decode.
- `DsmEventDecoder.psml` - The powershell module, which queries the event logs for information, calls the `FormatDsmEventLog` cmdlet to parse and decode the event information.

Decoded output

The information decoded into a CSV format consists of several sections as described below.

1. The first section describes the input arguments to the powershell decoder script.
2. The second section is a detailed dump of the BaseComponent and EventComponent XML files. You can use this section to manually decode the event data if the automated process runs into an error with the event data. This section is also useful if only the decoded results are provided to technical support rather than the original *.EVTX files.
3. The last section is the actual decoded events. Note that the entire event log is decoded, not just the event specific information. Furthermore, an attempt to decode the Microsoft MPIO-generated events is provided for convenience.

Limitations

The following items list the limitations for the decoding process.

- If a large number of records are present the decoding process might take some time.
- CSV format is currently the only supported output format.

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 the multipath software

If a disruption of one or more physical paths occurs, the multipath software maintains an active path to the underlying network storage. The multipath software presents the operating system with a single virtual device that represents the active physical paths to the storage and manages the failover process that updates the virtual device.

About this task

If you have not already done so, you can perform a full installation of SANtricity Storage Manager software. If you need to install only the utilities package and the failover package, you can perform a

custom installation. For more information regarding installation methods and customizations see [Configuration options](#) on page 6.

Steps

1. Download the SANtricity software build from [NetApp Support](#).
2. Execute the SANtricity installer. Double-click the SMIA*.exe install package to execute.
Select **Typical (Full) installation** or **Custom installation**, and choose to install the utilities package and the failover package.

Compatibility and migration

Operating systems supported

The DSM is supported on Windows Server 2008 R2 and later.

Storage interfaces supported

The DSM supports any protocol supported by MPIO, including Fiber Channel, SAS, and iSCSI.

SAN-Boot support

The DSM supports booting Windows from storage that is externally attached to the host.

Running the DSM in a hyper-v guest with pass-through disks

Consider a scenario where you map storage to a Windows Server 2008 R2 parent partition. You use the **Settings > SCSI Controller > Add Hard Drive** command to attach that storage as a pass-through disk to the SCSI controller of a Hyper-V guest running Windows Server 2008. By default, some SCSI commands are filtered by Hyper-V, so the DSM multipath driver fails to run properly.

To work around this issue, you must disable SCSI command filtering. Run the following PowerShell script in the parent partition to determine if SCSI pass-through filtering is enabled or disabled:

```
# Powershell Script: Get_SCSI_Passthrough.ps1
$TargetHost=$args[0] foreach ($Child in Get-WmiObject
-namespace root\virtualization Msvm_ComputerSystem
-Filter "ElementName='$TargetHost'") { $vmData=Get-WmiObject
-namespace root\virtualization -Query "Associators of {$Child}
Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"
Write-Host "Virtual Machine:" $vmData.ElementName
Write-Host "Currently Bypassing SCSI Filtering:"
$vmData.AllowFullSCSICommandSet
}
```

If necessary, run the following PowerShell script in the parent partition to disable SCSI Filtering:

```
# Powershell Script: Set_SCSI_Passthrough.ps1
$TargetHost=$args[0]
$vsManagementService=gwmi MSVM_VirtualSystemManagementService
-namespace "root\virtualization" for each ($Child in Get-WmiObject
-namespace root\virtualization Msvm_ComputerSystem
-Filter "ElementName='$TargetHost'") { $vmData=Get-WmiObject
-namespace root\virtualization -Query "Associators of {$Child}
Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"
$vmData.AllowFullSCSICommandSet=$true
$vsManagementService.ModifyVirtualSystem($Child,
$vmData.PSBase.GetText(1))|out-null
}
```

```
}
}
```

Installation and removal

Installing or updating DSM

About this task

Perform the steps in this task to install SANtricity Storage Manager and the DSM or to upgrade from an earlier release of SANtricity Storage Manager and the DSM on a system with a Windows operating system. For a clustered system, perform these steps on each node of the system, one node at a time.

Steps

1. Open the SANtricity Storage Manager SMIA installation program, which is available from your storage vendor's website.
2. Click **Next**.
3. Accept the terms of the license agreement, and click **Next**.
4. Select **Custom**, and click **Next**.
5. Select the applications that you want to install.
6. Click the name of an application to see its description.
7. Select the check box next to an application to install it.
8. Click **Next**.

If you have a previous version of the software installed, you receive a warning message: Existing versions of the following software already reside on this computer. If you choose to continue, the existing versions are overwritten with new versions.

9. If you receive this warning and want to update SANtricity Storage Manager, click **OK**.
10. Select whether to automatically start the Event Monitor. Click **Next**.
Start the Event Monitor for the one I/O host on which you want to receive alert notifications. Do not start the Event Monitor for all other I/O hosts attached to the storage array or for computers that you use to manage the storage array.
11. Click **Next**.
12. If you receive a warning about anti-virus or backup software that is installed, click **Continue**.
13. Read the pre-installation summary, and click **Install**.
14. Wait for the installation to complete, and click **Done**.

Uninstalling DSM

Reconfigure the connections between the host and the storage array to remove any redundant connections before you uninstall SANtricity Storage Manager and the DSM multipath driver.

About this task

Attention: To prevent loss of data, the host from which you are removing SANtricity Storage Manager and the DSM must have only one path to the storage array.

Steps

1. From the Windows Start menu, select **Control Panel**.
The Control Panel window appears.
2. In the Control Panel window, double-click **Add or Remove Programs**.
The Add or Remove Programs window appears.
3. Select **SANtricity Storage Manager**.
4. Click the **Remove** button to the right of the SANtricity Storage Manager entry.

Collecting trace events from a target machine

There are several utilities and tools that can be used to collect Trace Events. These tools and utilities typically establish a new trace session, along with specifying what flags and level of tracing to capture. When capturing is complete, the trace session is stopped and the capture buffers flushed of any cached information.

Control files

Several tools and utilities require knowing the GUID of the provider as well as trace flags and level. If you want only to collect information for a single provider, you can provide the GUID and trace settings through one or more command-line arguments. To capture from multiple sources, use Control Files. The Control File format is typically:

```
{GUID} [Flags Level]
```

For example:

```
C:>type mppdsmctl
{706a8802-097d-43C5-ad89-8863e84774c6} 0x0000FFFF 0xF
```

Logman

The Logman tool manages and schedules performance counter and event trace log collections on local and remote systems, and is provided in-box with each OS installation. There is no explicit requirement for the DSM Trace Provider to be registered before you can use Logman to capture trace events, although for end-user convenience the DSM should be registered during installation.

Viewing a list of available providers

To view a list of available providers:

```
C:>logman query providers
```

By default the DSM does not appear in this list unless it has previously been registered.

Establishing a new trace session

To establish a new trace session:

```
C:>logman create trace <session_name> -ets -nb 16 256 -bs 64 -o
<logfile> -pf <control_file>
```

Where:

- <session_name>: Name of the trace session (ex. "mppdsm")
- <control_file>: Trace control file.

Determine status of trace sessions

To determine whether a trace session is running, using the 'query' option. In this example an 'mppdsm' trace session has been created and shown as running:

```
C:\Users\Administrator>logman query -ets
```

Data Collector Set	Type	Status
AIEventLog	Trace	Running
Audio	Trace	Running
DiagLog	Trace	Running
EventLog-Application	Trace	Running
EventLog-System	Trace	Running
NtfsLog	Trace	Running
SQMLogger	Trace	Running
UAL_Usermode_Provider	Trace	Running
UBPM	Trace	Running
WdiContextLog	Trace	Running
umstartup	Trace	Running
Terminal-Services-Core	Trace	Running
Terminal-Services-RPC-Client	Trace	Running
Terminal-Services-Unified-APIs	Trace	Running
Terminal-Services-IP-Virtualization	Trace	Running
Terminal-Services-SessionEnv	Trace	Running
Terminal-Services-SessionMsg	Trace	Running
MSDTC_TRACE_SESSION	Trace	Running
UAL_Kernelmode_Provider	Trace	Running
mppdsm	Trace	Running
WBEEngine	Trace	Running

The command completed successfully.

82008-04

The following command can be used to get more detailed information about the trace session. In this example, the 'mppdsm' session is detailed:

```
C:\Users\Administrator>logman query mppdsm -ets
```

```
Name: mppdsm
Status: Running
Root Path: C:\Users\Administrator
Segment: Off
Schedules: On

Name: mppdsm\mppdsm
Type: Trace
Output Location: C:\Users\Administrator\dsm.log
Append: Off
Circular: Off
Overwrite: Off
Buffer Size: 64
Buffers Lost: 0
Buffers Written: 1
Buffer Flush Timer: 0
Clock Type: Performance
File Mode: File

Provider:
Name: {706A8802-097D-43C5-AD89-8863E84774C6}
Provider Guid: {706A8802-097D-43C5-AD89-8863E84774C6}
Level: 15
KeywordsAll: 0x0
KeywordsAny: 0xffff
Properties: 0
Filter Type: 0

The command completed successfully.
```

82008-05

Stopping a trace session

To stop a tracing session:

```
C:\Users\Administrator>logman stop <session_name> -ets  
The command completed successfully.
```

Deleting a trace session

To delete a tracing session:

```
C:\Users\Administrator>logman delete <session_name>  
The command completed successfully.
```

Enabling a boot-time trace session

Enabling boot-time tracing is done by appending "autosession" to the session name:

```
logman create trace "autosession\<session_name>"  
-o <logfile> -pf <control_file>
```

For example:

```
C:\Users\Administrator>logman create trace "autosession\mppdsm"  
-o mppdsmtrace.etl -pf mppdsm.ctl  
The command completed successfully.
```

Boot-Time sessions can be stopped and deleted just like any other session.

Note: You need to register the DSM as a provider with WMI or boot-time logging does not occur.

Disabling a boot-time trace session

To disable a boot-time trace session:

```
C:\Users\Administrator>logman delete "autosession\mppdsm"  
The command completed successfully.
```

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 Windows Express Guide*](#)

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

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.

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 notes, 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.

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.

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

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