

Lenovo

ThinkSystem® SAN OS 11.70.1 Installing and Configuring for Windows® Express Guide



First edition (September 2021)

© Copyright Lenovo 2021.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

Contents

Deciding whether to use this Express Guide	1
Lenovo Storage Interoperation Center (LSIC)	2
Understanding the workflow	3
Verifying the configuration is supported	4
Lenovo Storage Interoperation Center (LSIC)	4
Configuring management port IP addresses.....	5
Configuring the multipath software	6
Install ThinkSystem Windows DSM package	6
Hyper-V guest OS configuration	6
Performing FC-specific tasks	8
Determining host WWPNs and making the recommended settings - FC.....	8
Lenovo Storage Interoperation Center (LSIC)	8
Configuring the switches - FC.....	8
FC worksheet	9
Performing iSCSI-specific tasks.....	11
Configuring the switches - iSCSI	11
Configuring networking - iSCSI	11
Configuring array-side networking - iSCSI.....	11
Configuring host-side networking - iSCSI	14
Verifying IP network connections—iSCSI	15
iSCSI worksheet	16
Performing SAS-specific tasks.....	18
Determining SAS host identifiers	18
SAS worksheet	18
Installing ThinkSystem Storage Manager Host Context Agent (HCA)	20
Access ThinkSystem System Manager and use the Setup Wizard	21
Creating a volume.....	23

Create workloads	23
Create volumes.....	24
Step 1: Select host for a volume	24
Step 2: Select a workload for a volume.....	25
Step 3: Add or edit volumes	26
Step 4: Review volume configuration	30
Defining a host in ThinkSystem System Manager	31
Create host automatically	31
Create host manually.....	31
Create host cluster.....	33
Mapping a volume to a host	35
Discovering storage on the host.....	36
Configuring storage on the host	37
Verifying storage access on the host.....	39
Where to find additional information	40
Contacting Support	41
Notices	42
Trademarks.....	42

Deciding whether to use this Express Guide

The express method for installing your storage array and accessing ThinkSystem System Manager is appropriate for setting up a standalone Windows host to a DE Series system. It is designed to get the storage system up and running as quickly as possible with minimal decision points.

The express method includes the following steps:

1. Setting up one of the following communication environments:
 - Fibre Channel (FC)
 - iSCSI
 - SAS
2. Creating logical volumes on the storage array.
3. Making the volume LUNs available to the data host.

This guide is based on the following assumptions:

Component	Assumptions
Hardware	<ul style="list-style-type: none">• You have used the Installation and Setup Instructions included with the controller shelves to install the hardware.• You have connected cables between the optional drive shelves and the array controllers.• You have applied power to the storage array.• You have installed all other hardware (for example, management station, switches) and made the necessary connections.
Host	<ul style="list-style-type: none">• You have made a connection between the storage array and the data host.• You have installed the host operating system.• You are not using Windows as a virtualized guest.• You are not configuring the data (I/O attached) host to boot from SAN.
Storage management station	<ul style="list-style-type: none">• You are using a 1 Gbps or faster management network.• You are using a separate station for management rather than the data (I/O attached) host.• You are using out-of-band management, in which a storage management station sends commands to the storage array through the Ethernet connections to the controller.• You have attached the management station to the same subnet as the storage management ports.
IP addressing	<ul style="list-style-type: none">• You have installed and configured a DHCP server.• You have not yet made an Ethernet connection between the management station and the storage array.
Storage provisioning	<ul style="list-style-type: none">• You will not use shared volumes.• You will create pools rather than volume groups.
Protocol: FC	<ul style="list-style-type: none">• You have made all host-side FC connections and activated switch zoning.• You are using Lenovo-supported FC HBAs and switches.

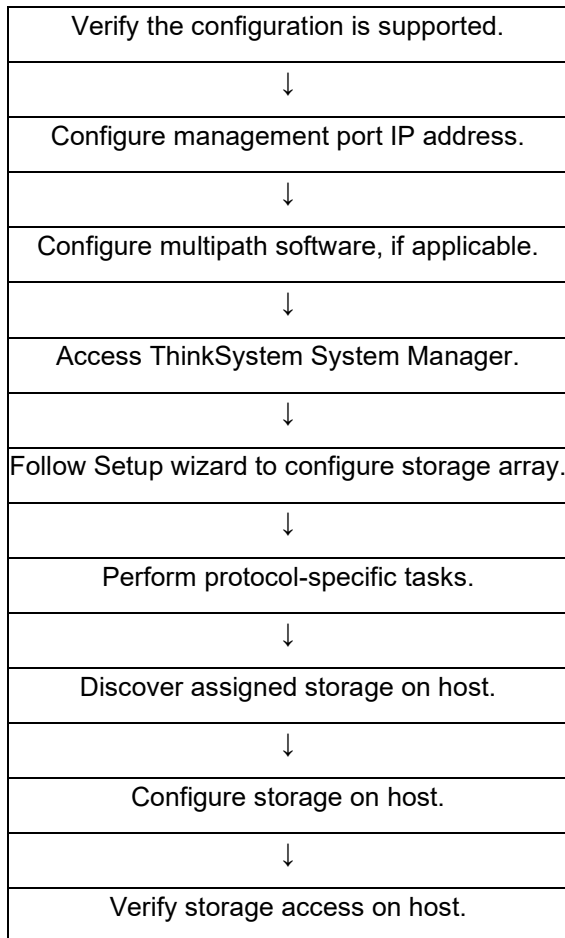
Component	Assumptions
	<ul style="list-style-type: none"> You are using FC HBA driver versions as listed in Interoperability Matrix on the DE Series Product Support Site.
Protocol: iSCSI	<ul style="list-style-type: none"> You are using Ethernet switches capable of transporting iSCSI traffic. You have configured the Ethernet switches according to the vendor's recommendation for iSCSI.
Protocol: SAS	<ul style="list-style-type: none"> You are using Lenovo-supported SAS HBAs. You are using SAS HBA driver versions as listed in Interoperability Matrix on the DE Series Product Support Site.

Related information

[Lenovo Storage Interoperation Center \(LSIC\)](#)

Understanding the workflow

This workflow guides you through the express method for configuring your storage array and ThinkSystem System Manager to make storage available to a host.



Verifying the configuration is supported

To ensure reliable operation, you create an implementation plan and then verify that the entire configuration is supported.

1. Go to the [DE Series Product Support Site](#).
2. Look for the *Interoperability Matrix* document, and click to download or view the file.

In this file, you may search for the product family that applies, as well as other criteria for the configuration such as Operating System, ThinkSystem SAN OS, and Host Multipath driver.

3. As necessary, make the updates for your operating system and protocol as listed in the table.

Operating system updates	Protocol	Protocol-related updates
You might need to install out-of-box drivers to ensure proper functionality and supportability. Each HBA vendor has specific methods for updating boot code and firmware. Refer to the support section of the vendor's website to obtain the instructions and software necessary to update the HBA boot code and firmware.	FC	Host bus adapter (HBA) driver, firmware, and bootcode
	iSCSI	Network interface card (NIC) driver, firmware and bootcode
	SAS	Host bus adapter (HBA) driver, firmware, and bootcode

Related information

[Lenovo Storage Interoperation Center \(LSIC\)](#)

Configuring management port IP addresses

In this express method for configuring communications between the management station and the storage array, you use Dynamic Host Configuration Protocol (DHCP) to provide IP addresses. Each controller has two storage management ports, and each management port will be assigned an IP address.

Before you begin

You have installed and configured a DHCP server on the same subnet as the storage management ports.

The following instructions refer to a storage array with two controllers (a duplex configuration).

1. If you have not already done so, connect an Ethernet cable to the management station and to management port 1 on each controller (A and B).

The DHCP server assigns an IP address to port 1 of each controller.

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

Important: If you disconnect and reconnect the Ethernet cable, or if the storage array is power-cycled, DHCP assigns IP addresses again. This process occurs until static IP addresses are configured. It is recommended that you avoid disconnecting the cable or power-cycling the array.

If the storage array cannot get DHCP-assigned IP addresses within 30 seconds, the following default IP addresses are set:

- Controller A, port 1: 169.254.128.101
 - Controller B, port 1: 169.254.128.102
 - Subnet mask: 255.255.0.0
2. Locate the MAC address label on the back of each controller, and then provide your network administrator with the MAC address for port 1 of each controller.

Your network administrator needs the MAC addresses to determine the IP address for each controller. You will need the IP addresses to connect to your storage system through your browser.

Configuring the multipath software

Multipath software provides a redundant path to the storage array in case one of the physical paths is disrupted. Before you can use multipathing, you need to enable Windows Multipathing feature and the ThinkSystem Windows DSM package. This package contains the multipath software for Windows.

Windows installations use the native MPIO Device Specific Module (DSM) driver for failover. When you install ThinkSystem Storage Manager, the DSM driver is installed and enabled.

When installing the DSM for windows, you must first install MPIO through the Server Manager console. Once completed you must install the DSM by installing the ThinkSystem Storage Manager package. This package contains the DSM, the context agent, and the SMCLI. While Storage Manager is a deprecated application, it is important that if you uninstall Storage Manager, you don't uninstall the DSM. MPIO must be installed prior to installing the DSM for MPIO to work properly.

Install ThinkSystem Windows DSM package

Perform the following steps to install the ThinkSystem Windows DSM package and use the multipath package for Windows.

Before you begin

- You must have the correct administrator or superuser privileges.
- Enable the Multipathing Feature under PowerShell with administrator privilege and reboot: `Add-WindowsFeature -Name 'Multipath-IO' -Restart`

You will install the ThinkSystem Windows DSM package on the management station.

1. Download the ThinkSystem Storage Manager package from [DE Series Product Support Site Drivers&Software > Software and Utilities](#)
2. Run the ThinkSystem Storage Manager. Double-click the install package to execute.
3. Use the installation wizard to install the package on the management station.

Hyper-V guest OS configuration

- **Linux guest OS under Hyper-V**

The `\"udev\"` facility is the best way to automatically configure disk timeouts when new disks are added to a VM and at boot time. A simple way to determine the VID /PID (Vendor ID/Product ID) of any disks currently added to the VM is to run the following command and look at the vendor and model information: `cat /proc/scsi/scsi`. In order to set this up, create the file `/etc/udev/rules.d/99-msft-udev.rules` with the following content:

```
ACTION=="add",
SUBSYSTEMS=="scsi",
ATTRS{vendor}=="LENOVO ",
ATTRS{model}=="DE_Series",
RUN+="/bin/sh -c 'echo 180 >/sys$DEVPATH/timeout'"
ACTION=="add",
```

```
SUBSYSTEMS=="scsi",  
ATTRS{vendor}=="Msft  ",  
ATTRS{model}=="Virtual Disk  ",  
RUN+="bin/sh -c 'echo 180 >/sys$DEVPATH/timeout'"
```

The vendor and model fields in the udev rules file have fixed-widths of 8 and 16 characters respectively. Note the additional white space padding in the example above. Once this file has been modified, you can either restart udev or reboot the VM to activate the changes.

- **Windows guest OS under Hyper-V**

In order to configure disk I/O timeouts for Windows Hyper-V environments, you must edit a single registry key:

```
[HKEY_LOCAL_MACHINE\\SYSTEM\\CurrentControlSet\\Services\\Disk]  
  
\\TimeOutValue\\=dword:000000b4
```

Once that's done, you must restart the VM for the changes to take effect.

For Linux and Windows guest OS under Hyper-V environment, you will need to make additional changes

Performing FC-specific tasks

For the Fibre Channel protocol, you configure the switches and determine the host port identifiers.

Determining host WWPNs and making the recommended settings - FC

You install an FC HBA utility so you can view the worldwide port name (WWPN) of each host port. Additionally, you can use the HBA utility to change any settings recommended in the *Interoperability Matrix* document for the supported configuration.

Guidelines for HBA utilities:

- Most HBA vendors offer an HBA utility. You will need the correct version of HBA for your host operating system and CPU. Examples of FC HBA utilities include:
 - Emulex OneCommand Manager for Emulex HBAs
 - QLogic QConverge Console for QLogic HBAs
 - Host I/O ports might automatically register if the host context agent is installed.
1. Download the appropriate utility from your HBA vendor's web site.
 2. Install the utility.
 3. Select the appropriate settings in the HBA utility.

Appropriate settings for your configuration are listed in the *Interoperability Matrix* document. Go to [DE Series Product Support Site](#), click on the **Knowledge Base & Guides** tab, and look for the *Interoperability Matrix* document.

Related information

[Lenovo Storage Interoperation Center \(LSIC\)](#)

Configuring the switches - FC

Configuring (zoning) the Fibre Channel (FC) switches enables the hosts to connect to the storage array and limits the number of paths. You zone the switches using the management interface for the switches.

Before you begin

- You must have administrator credentials for the switches.
- You must have used your HBA utility to discover the WWPN of each host initiator port and of each controller target port connected to the switch.

For details about zoning your switches, see the switch vendor's documentation.


You must zone by WWPN, not by physical port. Each initiator port must be in a separate zone with all of its corresponding target ports.

1. Log in to the FC switch administration program, and then select the zoning configuration option.
2. Create a new zone that includes the first host initiator port and that also includes all of the target ports that connect to the same FC switch as the initiator.
3. Create additional zones for each FC host initiator port in the switch.
4. Save the zones, and then activate the new zoning configuration.

FC worksheet

You can use this worksheet to record FC storage configuration information. You need this information to perform provisioning tasks.

The illustration shows a host connected to an DE Series storage array in two zones. One zone is indicated by the blue line; the other zone is indicated by the red line. Any single port has two paths to the storage (one to each controller).



Host identifiers

Callout No.	Host (initiator) port connections	WWPN
1	Host	<i>not applicable</i>
2	Host port 0 to FC switch zone 0	
7	Host port 1 to FC switch zone 1	

Target identifiers

Callout No.	Array controller (target) port connections	WWPN
3	Switch	<i>not applicable</i>
6	Array controller (target)	<i>not applicable</i>
5	Controller A, port 1 to FC switch 1	
9	Controller A, port 2 to FC switch 2	
4	Controller B, port 1 to FC switch 1	
8	Controller B, port 2 to FC switch 2	

Mapping host

Mapping host name	
Host OS type	

Performing iSCSI-specific tasks

For the iSCSI protocol, you configure the switches and configure networking on the array side and the host side. Then you verify the IP network connections.

Configuring the switches - iSCSI

You configure the switches according to the vendor's recommendations for iSCSI. These recommendations might include both configuration directives as well as code updates.

You must ensure the following:

- You have two separate networks for high availability. Make sure that you isolate your iSCSI traffic to separate network segments.
- You have enabled send and receive hardware flow control **end to end**.
- You have disabled priority flow control.
- If appropriate, you have enabled jumbo frames.

Note: Port channels/LACP is not supported on the controller's switch ports. Host-side LACP is not recommended; multipathing provides the same, and in some cases better, benefits.

Configuring networking - iSCSI

You can set up your iSCSI network in many ways, depending on your data storage requirements.

Consult your network administrator for tips on selecting the best configuration for your environment.

An effective strategy for configuring the iSCSI network with basic redundancy is to connect each host port and one port from each controller to separate switches and partition each set of host and controller ports on separate network segments using VLANs.

You must enable send and receive hardware flow control **end to end**. You must disable priority flow control.

If you are using jumbo frames within the IP SAN for performance reasons, make sure to configure the array, switches, and hosts to use jumbo frames. Consult your operating system and switch documentation for information on how to enable jumbo frames on the hosts and on the switches. To enable jumbo frames on the array, complete the steps in *Configuring array-side networking—iSCSI*.

Note: Many network switches have to be configured above 9,000 bytes for IP overhead. Consult your switch documentation for more information.

Configuring array-side networking - iSCSI

You use the ThinkSystem System Manager GUI to configure iSCSI networking on the array side.

Before you begin

- You must know the IP address or domain name for one of the storage array controllers.

- You or your system administrator must have set up a password for the System Manager GUI, or you must configure Role-Based Access Control (RBAC) or LDAP and a directory service for the appropriate security access to the storage array. See the *ThinkSystem System Manager online help* for more information about Access Management.

This task describes how to access the iSCSI port configuration from the Hardware page. You can also access the configuration from **System > Settings > Configure iSCSI ports**.

- From your browser, enter the following URL: `https://<DomainNameOrIPAddress>`

IPAddress is the address for one of the storage array controllers.

The first time ThinkSystem System Manager is opened on an array that has not been configured, the Set Administrator Password prompt appears. Role-based access management configures four local roles: **admin**, **support**, **security**, and **monitor**. The latter three roles have random passwords that cannot be guessed. After you set a password for the **admin** role you can change all of the passwords using the **admin** credentials. See *ThinkSystem System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the **admin** role in the **Set Administrator Password** and **Confirm Password** fields, and then select the **Set Password** button.

When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.

- Close the Setup wizard.

You will use the wizard later to complete additional setup tasks.

- Select **Hardware**.

- If the graphic shows the drives, click **Show back of shelf**.

The graphic changes to show the controllers instead of the drives.

- Click the controller with the iSCSI ports you want to configure.

The controller's context menu appears.

- Select **Configure iSCSI ports**.

The Configure iSCSI Ports dialog box opens.

- In the drop-down list, select the port you want to configure, and then click **Next**.

- Select the configuration port settings, and then click **Next**.

To see all port settings, click the **Show more port settings** link on the right of the dialog box.

Port Setting	Description
Configured Ethernet port speed	Select the desired speed.

Port Setting	Description
	<p>The options that appear in the drop-down list depend on the maximum speed that your network can support (for example, 10 Gbps).</p> <p>Note: The optional iSCSI host interface cards in the DE6000H and DE6000F controllers do not auto-negotiate speeds. You must set the speed for each port to either 10 Gb or 25 Gb. All ports must be set to the same speed.</p>
Enable IPv4 / Enable IPv6	Select one or both options to enable support for IPv4 and IPv6 networks.
TCP listening port (Available by clicking Show more port settings.)	<p>If necessary, enter a new port number.</p> <p>The listening port is the TCP port number that the controller uses to listen for iSCSI logins from host iSCSI initiators. The default listening port is 3260. You must enter 3260 or a value between 49152 and 65535.</p>
MTU size (Available by clicking Show more port settings.)	<p>If necessary, enter a new size in bytes for the Maximum Transmission Unit (MTU).</p> <p>The default Maximum Transmission Unit (MTU) size is 1500 bytes per frame. You must enter a value between 1500 and 9000.</p>
Enable ICMP PING responses	Select this option to enable the Internet Control Message Protocol (ICMP). The operating systems of networked computers use this protocol to send messages. These ICMP messages determine whether a host is reachable and how long it takes to get packets to and from that host.

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

10. Configure the IPv4 and/or IPv6 settings, either automatically or manually. To see all port settings, click the **Show more settings** link on the right of the dialog box.

Port setting	Description
Automatically obtain configuration	Select this option to obtain the configuration automatically.
Manually specify static configuration	Select this option, and then enter a static address in the fields. For IPv4, include the network subnet mask and gateway. For IPv6, include the routable IP address and router IP address.
Enable VLAN support	Important: This option is only available in an iSCSI environment.

Port setting	Description
(Available by clicking Show more settings .)	<p>It is not available in an NVMe over RoCE environment.</p> <p>Select this option to enable a VLAN and enter its ID. A VLAN is a logical network that behaves like it is physically separate from other physical and virtual local area networks (LANs) supported by the same switches, the same routers, or both.</p>
<p>Enable ethernet priority</p> <p>(Available by clicking Show more settings.)</p>	<p>Important: This option is only available in an iSCSI environment.</p> <p>It is not available in an NVMe over RoCE environment.</p> <p>Select this option to enable the parameter that determines the priority of accessing the network. Use the slider to select a priority between 1 and 7.</p> <p>In a shared local area network (LAN) environment, such as Ethernet, many stations might contend for access to the network. Access is on a first-come, first-served basis. Two stations might try to access the network at the same time, which causes both stations to back off and wait before trying again. This process is minimized for switched Ethernet, where only one station is connected to a switch port.</p>

11. Click **Finish**.

12. Close System Manager.

Configuring host-side networking - iSCSI

You must configure iSCSI networking on the host side so that the Microsoft iSCSI Initiator can establish sessions with the array.

Before you begin

- You have fully configured the switches that will be used to carry iSCSI storage traffic.
- You must have enabled send and receive hardware flow control **end to end** and disabled priority flow control.
- You have completed the array side iSCSI configuration.
- You must know the IP address of each port on the controller.

These instructions assume that two NIC ports will be used for iSCSI traffic.

1. Disable unused network adapter protocols.

These protocols include, but are not limited to, QoS, File and Print Sharing, and NetBIOS.

2. Execute > iscsicpl.exe from a terminal window on the host to open the **iSCSI Initiator Properties** dialog box.

3. On the **Discovery** tab, select **Discover Portal** , and then enter the IP address of one of the iSCSI target ports.
4. On the **Targets** tab, select the first target portal you discovered and then select **Connect**.
5. Select **Enable multi-path**, select **Add this connection to the list of Favorite Targets**, and then select **Advanced**.
6. For **Local adapter**, select **Microsoft iSCSI Initiator**.
7. For **Initiator IP**, select the IP address of a port on the same subnet or VLAN as one of the iSCSI targets.
8. For **Target IP**, select the IP address of a port on the same subnet as the **Initiator IP** selected in the step above.
9. Retain the default values for the remaining check boxes, and then select **OK**.
10. Select **OK** again as you return to the **Connect to Target** dialog box.
11. Repeat this procedure for each initiator port and session (logical path) to the storage array that you want to establish.

Verifying IP network connections—iSCSI

You verify Internet Protocol (IP) network connections by using ping tests to ensure the host and array are able to communicate.

1. Select **Start > All Programs > Accessories > Command Prompt**, and use the Windows CLI to run one of the following commands, depending on whether jumbo frames are enabled:

- If jumbo frames are not enabled, run this command:

```
ping -s <hostIP> <targetIP>
```

- If jumbo frames are enabled, run the `ping` command with a payload size of 8,972 bytes. The IP and ICMP combined headers are 28 bytes, which when added to the payload, equals 9,000 bytes. The `-f` switch sets the don't fragment (DF) bit. The `-l` switch allows you to set the size. These options allow jumbo frames of 9,000 bytes to be successfully transmitted between the iSCSI initiator and the target.

```
ping -l 8972 -f <iSCSI_target_IP_address>
```

In this example, the iSCSI target IP address is 192.0.2.8.

```
C:\>ping -l 8972 -f 192.0.2.8

Pinging 192.0.2.8 with 8972 bytes of data:

Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64
Reply from 192.0.2.8: bytes=8972 time=2ms TTL=64
```

Ping statistics for 192.0.2.8:

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

Approximate round trip times in milli-seconds:

Minimum = 2ms, Maximum = 2ms, Average = 2ms

- Issue a `ping` command from each host's initiator address (the IP address of the host Ethernet port used for iSCSI) to each controller iSCSI port. Perform this action from each host server in the configuration, changing the IP addresses as necessary.


Note: If the command fails (for example, returns "Packet needs to be fragmented but DF set"), verify the MTU size (jumbo frame support) for the Ethernet interfaces on the host server, storage controller, and switch ports.

iSCSI worksheet

You can use this worksheet to record iSCSI storage configuration information. You need this information to perform provisioning tasks.

Recommended configuration

Recommended configurations consist of two initiator ports and four target ports with one or more VLANs.



Target IQN

Callout No.	Target port connection	IQN
2	Target port	

Mappings host name

Callout No.	Host information	Name and type
1	Mappings host name	
	Host OS type	

Performing SAS-specific tasks

Determining SAS host identifiers

For the SAS protocol, you find the SAS addresses using the HBA utility, then use the HBA BIOS to make the appropriate configuration settings.


Guidelines for HBA utilities:

- Most HBA vendors offer an HBA utility. Depending on your host operating system and CPU, use either the LSI-sas2flash(6G) or sas3flash(12G) utility.
 - Host I/O ports might automatically register if the host context agent is installed.
1. Download the LSI-sas2flash(6G) or sas3flash(12G) utility from your HBA vendor's web site.
 2. Install the utility.
 3. Use the HBA BIOS to select the appropriate settings for your configuration.

Go to [DE Series Product Support Site](#), click on the **Knowledge Base & Guides** tab, and look for the *Interoperability Matrix* document for recommendations.

SAS worksheet

You can use this worksheet to record SAS storage configuration information. You need this information to perform provisioning tasks.



Host Identifiers

Callout No.	Host (initiator) port connections	SAS address
1	Host	<i>not applicable</i>
2	Host (initiator) port 1 connected to Controller A, port 1	
3	Host (initiator) port 1 connected to Controller B, port 1	
4	Host (initiator) port 2 connected to Controller A, port 1	

Callout No.	Host (initiator) port connections	SAS address
5	Host (initiator) port 2 connected to Controller B, port 1	

Target Identifiers

Recommended configurations consist of two target ports.

Mappings Host

Mappings Host Name	
Host OS Type	

Installing ThinkSystem Storage Manager Host Context Agent (HCA)

When you install the ThinkSystem Storage Manager software on your management station, the Host Context Agent is installed that helps the host push configuration information to the storage array controllers through the I/O path. For the 11.70.1R1 or later release, Storage Manager (Host Utilities) can only be installed on host servers. All multi-system management functions have been moved to SAN Manager.

1. Download the ThinkSystem Storage Manager package from [DE Series Product Support Site Drivers&Software > Software and Utilities](#).
2. Run the ThinkSystem Storage Manager. Double-click the install package to execute.
3. Use the installation wizard to install the package on the management station.

Access ThinkSystem System Manager and use the Setup Wizard

You use the Setup wizard in ThinkSystem System Manager to configure your storage array.

Before you begin

- You have ensured that the device from which you will access ThinkSystem System Manager contains one of the following browsers:

Browser	Minimum version
Google Chrome	47
Microsoft Internet Explorer	11
Microsoft Edge	EdgeHTML 12
Mozilla Firefox	31
Safari	9

- You are using out-of-band management.

If you are an iSCSI user, make sure you have closed the Setup wizard while configuring iSCSI.

The wizard automatically relaunches when you open System Manager or refresh your browser and *at least one* of the following conditions is met:

- No pools and volume groups are detected.
- No workloads are detected.
- No notifications are configured.

If the Setup wizard does not automatically appear, contact technical support.

- From your browser, enter the following URL: `https://<DomainNameOrIPAddress>`

IPAddress is the address for one of the storage array controllers.

The first time ThinkSystem System Manager is opened on an array that has not been configured, the Set Administrator Password prompt appears. Role-based access management configures four local roles: **admin**, **support**, **security**, and **monitor**. The latter three roles have random passwords that cannot be guessed. After you set a password for the **admin** role you can change all of the passwords using the **admin** credentials. See *ThinkSystem System Manager online help* for more information on the four local user roles.

- Enter the System Manager password for the **admin** role in the **Set Administrator Password** and **Confirm Password** fields, and then select the **Set Password** button.

When you open System Manager and no pools, volumes groups, workloads, or notifications have been configured, the Setup wizard launches.

- Use the Setup wizard to perform the following tasks:

- Verify hardware (controllers and drives)** – Verify the number of controllers and drives in the storage array. Assign a name to the array.

- **Verify hosts and operating systems** – Verify the host and operating system types that the storage array can access.
 - **Accept pools** – Accept the recommended pool configuration for the express installation method. A pool is a logical group of drives.
 - **Configure alerts** – Allow System Manager to receive automatic notifications when a problem occurs with the storage array.
 - **Enable AutoSupport** – Automatically monitor the health of your storage array and have dispatches sent to technical support.
4. If you have not already created a volume, create one by going to **Storage > Volumes > Create > Volume**.

For more information, see the online help for ThinkSystem System Manager.

Creating a volume

You create storage by first creating a workload for a specific application type. Next, you add storage capacity to the workload by creating volumes with similar underlying volume characteristics.

Create workloads

You can create workloads for any type of application.

About this task

A workload is a storage object that supports an application. You can define one or more workloads, or instances, per application. For some applications, the system configures the workload to contain volumes with similar underlying volume characteristics. These volume characteristics are optimized based on the type of application the workload supports. For example, if you create a workload that supports a Microsoft SQL Server application and then subsequently create volumes for that workload, the underlying volume characteristics are optimized to support Microsoft SQL Server.

System Manager recommends an optimized volume configuration only for the following application types:

- Microsoft® SQL Server™
- Microsoft® Exchange Server™
- Video Surveillance applications
- VMware ESXi™ (for volumes to be used with Virtual Machine File System)

Keep these guidelines in mind:

- *When using an application-specific workload*, the system recommends an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. You can review the recommended volume configuration, and then edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- *When using other application types*, you manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

1. Select **Storage > Volumes**.
2. Select **Create > Workload**.

The **Create Application Workload** dialog box appears.

3. Use the drop-down list to select the type of application that you want to create the workload for and then type a workload name.
4. Click **Create**.

You are ready to add storage capacity to the workload you created. Use the **Create Volume** option to create one or more volumes for an application, and to allocate specific amounts of capacity to each volume.

Create volumes

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

Most application types default to a user-defined volume configuration. Some application types have a smart configuration applied at volume creation. For example, if you are creating volumes for Microsoft Exchange application, you are asked how many mailboxes you need, what your average mailbox capacity requirements are, and how many copies of the database you want. System Manager uses this information to create an optimal volume configuration for you, which can be edited as needed.

The process to create a volume is a multi-step procedure.

Note: If you want to mirror a volume, first create the volumes that you want to mirror, and then use the **Storage > Volumes > Copy Services > Mirror a volume asynchronously** option.

1. [Step 1: Select host for a volume](#)

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

2. [Step 2: Select a workload for a volume](#)

Select a workload to customize the storage array configuration for a specific application, such as Microsoft SQL Server, Microsoft Exchange, Video Surveillance applications, or VMware. You can select "Other application" if the application you intend to use on this storage array is not listed.

3. [Step 3: Add or edit volumes](#)

System Manager may suggest a volume configuration based on the application or workload you selected. This volume configuration is optimized based on the type of application the workload supports. You can accept the recommended volume configuration, or you can edit it as needed. If you selected one of the "Other" applications, you must manually specify the volumes and characteristics you want to create.

4. [Step 4: Review volume configuration](#)

Review a summary of the volumes you intend to create and make any necessary changes.

Step 1: Select host for a volume

You create volumes to add storage capacity to an application-specific workload, and to make the created volumes visible to a specific host or host cluster. In addition, the volume creation sequence provides options to allocate specific amounts of capacity to each volume you want to create.

- Valid hosts or host clusters exist under the **Hosts** tile.
- Host port identifiers have been defined for the host.
- Before creating a DA-enabled volume, the host connection you are planning to use must support DA. If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes. ThinkSystem DE Series storage only supports DA between the controller and the drives.

Keep these guidelines in mind when you assign volumes:

- A host's operating system can have specific limits on how many volumes the host can access. Keep this limitation in mind when you create volumes for use by a particular host.

- You can define one assignment for each volume in the storage array.
- Assigned volumes are shared between controllers in the storage array.
- The same logical unit number (LUN) cannot be used twice by a host or a host cluster to access a volume. You must use a unique LUN.
- If you want to speed the process for creating volumes, you can skip the host assignment step so that newly created volumes are initialized offline.

Note: Assigning a volume to a host will fail if you try to assign a volume to a host cluster that conflicts with an established assignment for a host in the host clusters.

1. Select **Storage > Volumes**.
2. Select **Create > Volume**.

The **Create Volumes** dialog box appears.

3. From the drop-down list, select a specific host or host cluster to which you want to assign volumes, or choose to assign the host or host cluster at a later time.
4. To continue the volume creation sequence for the selected host or host cluster, click **Next**, and go to [Step 2: Select a workload for a volume](#).

The **Select Workload** dialog box appears

Step 2: Select a workload for a volume

Select a workload to customize the storage array configuration for a specific application, such as Microsoft SQL Server, Microsoft Exchange, Video Surveillance applications, or VMware. You can select "Other application" if the application you intend to use on this storage array is not listed.

This task describes how to create volumes for an existing workload.

- *When you are creating volumes using an application-specific workload*, the system may recommend an optimized volume configuration to minimize contention between application workload I/O and other traffic from your application instance. You can review the recommended volume configuration and edit, add, or delete the system-recommended volumes and characteristics using the **Add/Edit Volumes** dialog box.
- *When you are creating volumes using "Other" applications* (or applications without specific volume creation support), you manually specify the volume configuration using the **Add/Edit Volumes** dialog box.

1. Do one of the following:
 - Select the **Create volumes for an existing workload** option to create volumes for an existing workload.
 - Select the **Create a new workload** option to define a new workload for a supported application or for "Other" applications.
 - From the drop-down list, select the name of the application you want to create the new workload for.

Select one of the "Other" entries if the application you intend to use on this storage array is not listed.

- Enter a name for the workload you want to create.
2. Click **Next**.
 3. If your workload is associated with a supported application type, enter the information requested; otherwise, go to [Step 3: Add or edit volumes](#).

Step 3: Add or edit volumes

System Manager may suggest a volume configuration based on the application or workload you selected. This volume configuration is optimized based on the type of application the workload supports. You can accept the recommended volume configuration, or you can edit it as needed. If you selected one of the "Other" applications, you must manually specify the volumes and characteristics you want to create.

Before you begin

- The pools or volume groups must have sufficient free capacity.
- The maximum number of volumes allowed in a volume group is 256.
- The maximum number of volumes allowed in a pool depends on the storage system model:
 - 2,048 volumes (DE6000H, DE6000F series)
 - 512 volumes (DE2000H, DE4000H, DE4000F series)
- To create a Data Assurance (DA)-enabled volume, the host connection you are planning to use must support DA.

Selecting a DA capable pool or volume group

If you want to create a DA-enabled volume, select a pool or volume group that is DA capable (look for **Yes** next to "DA" in the pool and volume group candidates table).

DA capabilities are presented at the pool and volume group level in System Manager. DA protection checks for and corrects errors that might occur as data is transferred through the controllers down to the drives. Selecting a DA-capable pool or volume group for the new volume ensures that any errors are detected and corrected.

If any of the host connections on the controllers in your storage array do not support DA, the associated hosts cannot access data on DA-enabled volumes. ThinkSystem DE Series only supports DA between the controller and the drives.

- To create a secure-enabled volume, a security key must be created for the storage array.

Selecting a secure-capable pool or volume group

If you want to create a secure-enabled volume, select a pool or volume group that is secure capable (look for **Yes** next to "Secure-capable" in the pool and volume group candidates table).

Drive security capabilities are presented at the pool and volume group level in System Manager. Secure-capable drives prevent unauthorized access to the data on a drive that is physically removed from the storage array. A secure-enabled drive encrypts data during writes and decrypts data during reads using a unique *encryption key*.

A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.

About this task

You create volumes from pools or volume groups. The **Add/Edit Volumes** dialog box shows all eligible pools and volume groups on the storage array. For each eligible pool and volume group, the number of drives available and the total free capacity appears.

For some application-specific workloads, each eligible pool or volume group shows the proposed capacity based on the suggested volume configuration and shows the remaining free capacity in GiB. For other workloads, the proposed capacity appears as you add volumes to a pool or volume group and specify the reported capacity.

1. Choose one of these actions based on whether you selected Other or an application-specific workload:
 - **Other** – Click **Add new volume** in each pool or volume group that you want to use to create one or more volumes.

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	<p>Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes, the minimum capacity is 1 MiB, and the maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions – System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes – You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be</p>

Field	Description
	<p>impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size – The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> ▪ The I/O load from the host ▪ The modification priority of the volume ▪ The number of drives in the volume group ▪ The number of drive channels ▪ The processing power of the storage array controllers <p>When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>
Secure-capable	<p>Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.</p> <p>Drive Security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the Drive Security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the controllers and drives on a Storage Array.</p>

- **Application-specific workload** – Either click **Next** to accept the system-recommended volumes and characteristics for the selected workload, or click **Edit Volumes** to change, add, or delete the system-recommended volumes and characteristics for the selected workload.

Field	Description
Volume Name	A volume is assigned a default name by System Manager during the volume creation sequence. You can either accept the default name or provide a more descriptive one indicating the type of data stored in the volume.
Reported Capacity	Define the capacity of the new volume and the capacity units to use (MiB, GiB, or TiB). For Thick volumes , the minimum capacity is 1 MiB, and the

Field	Description
	<p>maximum capacity is determined by the number and capacity of the drives in the pool or volume group.</p> <p>Keep in mind that storage capacity is also required for copy services (snapshot images, snapshot volumes, volume copies, and remote mirrors); therefore, do not allocate all of the capacity to standard volumes.</p> <p>Capacity in a pool is allocated in 4-GiB increments. Any capacity that is not a multiple of 4 GiB is allocated but not usable. To make sure that the entire capacity is usable, specify the capacity in 4-GiB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the volume.</p>
Volume Type	Volume type indicates the type of volume that was created for an application-specific workload.
Segment Size	<p>Shows the setting for segment sizing, which only appears for volumes in a volume group. You can change the segment size to optimize performance.</p> <p>Allowed segment size transitions – System Manager determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the drop-down list. Allowed transitions usually are double or half of the current segment size. For example, if the current volume segment size is 32 KiB, a new volume segment size of either 16 KiB or 64 KiB is allowed.</p> <p>SSD Cache-enabled volumes – You can specify a 4-KiB segment size for SSD Cache-enabled volumes. Make sure you select the 4-KiB segment size only for SSD Cache-enabled volumes that handle small-block I/O operations (for example, 16 KiB I/O block sizes or smaller). Performance might be impacted if you select 4 KiB as the segment size for SSD Cache-enabled volumes that handle large block sequential operations.</p> <p>Amount of time to change segment size – The amount of time to change a volume's segment size depends on these variables:</p> <ul style="list-style-type: none"> ▪ The I/O load from the host ▪ The modification priority of the volume ▪ The number of drives in the volume group ▪ The number of drive channels ▪ The processing power of the storage array controllers <p>When you change the segment size for a volume, I/O performance is affected, but your data remains available.</p>
Secure-capable	Yes appears next to "Secure-capable" only if the drives in the pool or volume group are secure-capable.

Field	Description
	<p>Drive security prevents unauthorized access to the data on a drive that is physically removed from the storage array. This option is available only when the drive security feature has been enabled, and a security key is set up for the storage array.</p> <p>A pool or volume group can contain both secure-capable and non-secure-capable drives, but all drives must be secure-capable to use their encryption capabilities.</p>
DA	<p>Yes appears next to "DA" only if the drives in the pool or volume group support Data Assurance (DA).</p> <p>DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the controllers and drives on a Storage Array.</p>

2. To continue the volume creation sequence for the selected application, click **Next**, and go to [Step 4: Review volume configuration](#).

Step 4: Review volume configuration

Review a summary of the volumes that you intend to create and make any necessary changes.

1. Review the volumes you want to create. Click **Back** to make any changes.
2. When you are satisfied with your volume configuration, click **Finish**.

System Manager creates the new volumes in the selected pools and volume groups, and then displays the new volumes in the **All Volumes** table.

- Perform any operating system modifications necessary on the application host so that the applications can use the volume.
- Run either the host-based hot_add utility or an operating system-specific utility (available from a third-party vendor), and then run the SMdevices utility to correlate volume names with host storage array names.

The hot_add utility and the SMdevices utility are included as part of the SMutils package. The SMutils package is a collection of utilities to verify what the host sees from the storage array. It is included as part of the Storage Manager software installation.

Defining a host in ThinkSystem System Manager

You can create a host automatically or manually. To make it easier to give multiple hosts access to the same volumes, you can also create a host cluster.

Create host automatically

You can allow the Host Context Agent (HCA) to automatically detect the hosts, and then verify that the information is correct. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

Before you begin

The Host Context Agent (HCA) is installed and running on every host connected to the storage array. Hosts with the HCA installed and connected to the storage array are created automatically. To install the HCA, install ThinkSystem Storage Manager on the host and select the Host option. The HCA is not available on all supported operating systems. If it is not available, you must create the host manually.

1. Select **Storage > Hosts**.

The table lists the automatically created hosts.

2. Verify that the information provided by the HCA is correct (name, host type, host port identifiers).

If you need to change any of the information, select the host, and then click **View/Edit Settings**.

3. (Optional) If you want the automatically created host to be in a cluster, create a host cluster and add the host or hosts.

What happens next?

After a host is created automatically, the system displays the following items in the Hosts tile table:

- The host name derived from the system name of the host.
- The host identifier ports that are associated with the host.
- The Host Operating System Type of the host.

Create host manually

For hosts that cannot be automatically discovered, you can manually create a host. Creating a host is one of the steps required to let the storage array know which hosts are attached to it and to allow I/O access to the volumes.

About this task

Keep these guidelines in mind when you create a host:

- You must define the host identifier ports that are associated with the host.
- Make sure that you provide the same name as the host's assigned system name.
- This operation does not succeed if the name you choose is already in use.

- The length of the name cannot exceed 30 characters.

1. Select **Storage > Hosts**.
2. Click **Create > Host**.

The **Create Host** dialog box appears.

3. Select the settings for the host as appropriate.

Setting	Description
Name	Type a name for the new host.
Host operating system type	Select the operating system that is running on the new host from the drop-down list.
Host interface type	(Optional) If you have more than one type of host interface supported on your storage array, select the host interface type that you want to use.
Host ports	<p>Do one of the following:</p> <ul style="list-style-type: none"> • Select I/O Interface <p>Generally, the host ports should have logged in and be available from the drop-down list. You can select the host port identifiers from the list.</p> <ul style="list-style-type: none"> • Manual add <p>If a host port identifier is not displayed in the list, it means that the host port has not logged in. An HBA utility or the iSCSI initiator utility may be used to find the host port identifiers and associate them with the host.</p> <p>You can manually enter the host port identifiers or copy/paste them from the utility (one at a time) into the Host ports field.</p> <p>You must select one host port identifier at a time to associate it with the host, but you can continue to select as many identifiers that are associated with the host. Each identifier is displayed in the Host ports field. If necessary, you also can remove an identifier by selecting the X next to it.</p>
CHAP initiator	<p>(Optional) If you selected or manually entered a host port with an iSCSI IQN, and if you want to require a host that tries to access the storage array to authenticate using Challenge Handshake Authentication Protocol (CHAP), select the CHAP initiator checkbox. For each iSCSI host port you selected or manually entered, do the following:</p> <ul style="list-style-type: none"> • Enter the same CHAP secret that was set on each iSCSI host initiator for CHAP authentication. If you are using mutual CHAP authentication (two-way authentication that enables a host to validate itself to the storage array and for a storage array to validate itself to the host), you also must set the CHAP secret for the storage array at initial setup or by changing settings. • Leave the field blank if you do not require host authentication.

Setting	Description
	Currently, the only iSCSI authentication method used by System Manager is CHAP.

4. Click **Create**.

What happens next?

After the host is successfully created, the system creates a default name for each host port configured for the host (user label).

The default alias is < *Hostname_Port Number* >. For example, the default alias for the first port created for host IPT is IPT_1.

Create host cluster

You create a host cluster when two or more hosts require I/O access to the same volumes.

About this task

Keep these guidelines in mind when you create a host cluster:

- This operation does not start unless there are two or more hosts available to create the cluster.
- Hosts in host clusters can have different operating systems (heterogeneous).
- NVMe hosts in host clusters cannot be mixed with non-NVMe hosts.
- This operation does not succeed if the name you choose is already in use.
- The length of the name cannot exceed 30 characters.

1. Select **Storage > Hosts**.

2. Select **Create > Host Cluster**.

The **Create Host Cluster** dialog box appears.

3. Select the settings for the host cluster as appropriate.

Setting	Description
Name	Type the name for the new host cluster.
Select hosts to share volume access	Select two or more hosts from the drop-down list. Only those hosts that are not already part of a host cluster appear in the list.

4. Click **Create**.

If the selected hosts are attached to interface types that have different Data Assurance (DA) capabilities, a dialog appears with the message that DA will be unavailable on the host cluster. This unavailability prevents DA-enabled volumes from being added to the host cluster. Select **Yes** to continue or **No** to cancel.

DA increases data integrity across the entire storage system. DA enables the storage array to check for errors that might occur when data is moved between the controllers and the drives. Using DA for the new volume ensures that any errors are detected.

What happens next?

The new host cluster appears in the table with the assigned hosts in the rows beneath.

Mapping a volume to a host

For a host or host cluster to send I/O to a volume, you must assign the volume to the host or host cluster.

You can select a host or host cluster when you create a volume, or you can assign a volume to a host or host cluster later. A host cluster is a group of hosts. You create a host cluster to make it easy to assign the same volumes to multiple hosts.

Assigning volumes to hosts is flexible, allowing you to meet your particular storage needs.

- **Stand-alone host, not part of a host cluster** – You can assign a volume to an individual host. The volume can be accessed only by the one host.
- **Host cluster** – You can assign a volume to a host cluster. The volume can be accessed by all the hosts in the host cluster.
- **Host within a host cluster** – You can assign a volume to an individual host that is part of a host cluster. Even though the host is part of a host cluster, the volume can be accessed only by the individual host and not by any other hosts in the host cluster.

When volumes are created, logical unit numbers (LUNs) are assigned automatically. The LUN serves as the "address" between the host and the controller during I/O operations. You can change LUNs after the volume is created.

Discovering storage on the host

LUNs on your storage system appear as disks to the Windows host. When you add new LUNs, you must manually rescan the associated disks to discover them. The host does not automatically discover new LUNs.

Before you begin

You must be logged on as an administrator.

Steps

1. To discover the storage, run the following command from a Windows command prompt.

```
# echo rescan | diskpart
```

2. To verify the addition of new storage, run the following command.

```
# echo list disk | diskpart
```


Configuring storage on the host

A new LUN is offline and has no partition or file system when a Windows host first discovers it. You must bring the volume online and initialize it in Windows. Optionally, you can format the LUN with a file system.

Before you begin

The host must have discovered the LUN.

About this task

You can initialize the disk as a basic disk with a GPT or MBR partition table. Typically, you format the LUN with a file system such as New Technology File System (NTFS).

Steps

1. From a Windows command prompt, enter the diskpart context.

Example

```
> diskpart
```

2. View the list of available disks.

Example

```
> list disk
```

3. Select the disk to bring online.

Example

```
> select disk 1
```

4. Bring the disk online.

Example

```
> online disk
```

5. Create a partition.

Example

```
> create partition primary
```

Note: In Windows Server 2008 and later, you are prompted immediately after creating the partition to format the disk and give it a name. Select **Cancel** on the prompt to continue using these instructions for formatting and naming the partition.

6. Assign a drive letter.

Example

```
> assign letter=f
```

7. Format the disk.

Example

```
> format FS=NTFS LABEL="New Volume" QUICK
```

8. Exit the diskpart context.

Example

```
> exit
```

Verifying storage access on the host

Before using the volume, you verify that the host can write data to the LUN and read it back.

Before you begin

You must have initialized the LUN and formatted it with a file system.

1. Create and write to a file on the new LUN.

```
> echo test file > f:\test.txt
```

2. Read the file and verify data was written.

```
> type f:\test.txt
```

3. To verify that multipath is working, change the volume ownership.

- a. From the ThinkSystem System Manager GUI, go to **Storage > Volumes**, and then select **More > Change ownership**.
- b. On the **Change Volume Ownership** dialog box, use the **Preferred Owner** pull-down to select the other controller for one of the volumes in the list, and then confirm the operation.
- c. Verify that you can still access the files on the LUN.

```
> dir f:\
```

4. Find the target ID.

Note: The dsmUtil utility is case sensitive.

```
> C:\Program Files (x86)\DSMDrivers\mppdsm\dsmUtil.exe -a
```

5. View the paths to the LUN and verify that you have the expected number of paths. In the *<target ID>* portion of the command, use the target ID that you found in the previous step.

```
> C:\Program Files (x86)\DSMDrivers\mppdsm\dsmUtil.exe -g <target ID>
```

Where to find additional information

Use the resources listed here if you need additional information.

- ThinkSystem Storage DE Series on-line publications:
 - Hardware Installation and Maintenance Guide, Version 11.70.1
 - SAN Manager software, Version 5.1
 - System Manager software, Version 11.70.1
 - Embedded Command Line Interface, Version 11.70.1
- Lenovo Storage Interoperation Center
- Lenovo Press DE Series Storage

Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to

<https://datacentersupport.lenovo.com/serviceprovider> and use filter searching for different countries.

For Lenovo support telephone numbers, see <https://datacentersupport.lenovo.com/supportphonenumberlist> for your region support details.

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

Lenovo (United States), Inc. 8001 Development Drive

Morrisville, NC 27560 U.S.A.

Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT,

MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

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

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

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

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

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

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2021 Lenovo