


IBM Storage Host Attachment Kit
Version 2.6.0

User Guide



Note

Before using this document and the product it supports, read the information in "Notices" on page 105.

Edition notice

Publication number: GA32-1060-13. This publication applies to version 2.6.0 of the IBM Storage Host Attachment Kit and to all subsequent releases and modifications until otherwise indicated in a newer publication.

© Copyright IBM Corporation 2009, 2016.

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

Contents

Figures	vii
Tables	ix
About this guide.	xi
Who should use this guide	xi
Conventions used in this guide	xi
Related information and publications.	xii
Getting information, help, and service	xiii
IBM Publications Center	xiii
Sending your comments	xiii
Chapter 1. Introduction	1
Why use the host attachment kit?	1
Connectivity types	1
Concept diagram	1
Compatibility and requirements	2
The workflow	2
Download site	3
Before you proceed	3
Chapter 2. Preparation	5
Stage 1: Physically connecting the host to the IBM storage system	5
Option 1: Using Fibre Channel (FC) connections	6
Option 2: Using iSCSI connections	7
Stage 2: Setting connectivity parameters	8
Overview of Fibre Channel zoning	9
Identifying and checking Fibre Channel ports on the IBM storage system.	10
Overview of iSCSI settings	15
Using CHAP authentication for iSCSI connectivity	21
Stage 3: Preparing the host	22
Validating the system requirements	22
Validating your admin user privileges	23
Optional: booting the operating system from a storage area network (SAN)	23
Preparing AIX hosts (if applicable)	23
Chapter 3. Software installation	25
Option 1: Performing a clean local installation.	26
Installing on AIX, HP-UX, Linux, or Solaris.	26
Installing on Windows Server	28
Option 2: Upgrading from a previous local installation	31
Option 3: Using the portable package without installing	32
Using a mounted network drive	32
Using a portable USB flash drive	33
Removing the Host Attachment Kit software	33
Detaching the host and uninstalling the software.	33
Uninstalling the software without detaching the host	35
Displaying CLI manual pages	35
Chapter 4. Single host attachment	37
Stage 1: Using the CLI attachment utility (xiv_attach)	37
Issuing the xiv_attach command to connect to new targets	43
Stage 2: Defining the host on the storage system (if applicable)	43
Defining a host from the XIV GUI.	44

Defining a host from the CLI	45
Adding ports to a host from the XIV GUI	46
Defining a host from the IBM Hyper-Scale Manager	46
Stage 3: Mapping storage volumes to the host	49
Mapping storage volumes to the host from the XIV GUI	49
Mapping storage volumes to the host from the IBM Hyper-Scale Manager	50
Chapter 5. Cluster host attachment	53
Cluster installation prerequisites	54
Stage 1: Defining the cluster on the storage system	54
Defining a cluster from the XIV GUI	55
Defining a cluster from the CLI	57
Defining a cluster from the IBM Hyper-Scale Manager	57
Stage 2: Mapping storage volumes to the cluster	59
Mapping shared volumes to all the hosts in a cluster from the XIV GUI	60
Mapping private volumes to a specific host in a cluster from the XIV GUI	61
Mapping storage volumes to the cluster from the IBM Hyper-Scale Manager	62
Stage 3: Running <code>xiv_attach</code> on each cluster node.	64
Chapter 6. Additional tasks	65
Displaying information about connected storage systems	66
Displaying volume and additional information	70
Displaying volumes and other storage devices.	70
Displaying additional information	71
Setting Fibre Channel connectivity parameters.	73
Setting iSCSI connectivity parameters.	75
Displaying information about host configuration changes	79
Displaying information for <code>xiv_fc_admin</code> or <code>xiv_iscsi_admin</code>	79
Displaying information when working with <code>xiv_attach</code>	80
Chapter 7. Diagnostics	83
Using the host diagnostics utility	83
Using the Host Profiler utility	85
Chapter 8. Best practices.	89
Recommended operating system settings	89
Monitoring host connectivity	89
Monitoring from the XIV GUI	89
Monitoring from the CLI	90
Monitoring from the IBM Hyper-Scale Manager	90
Using the <code>rescan (-R)</code> option after new volumes are mapped	92
Scheduling the Host Profiler utility for weekly operation	92
Using the host-side acceleration (HSA)	93
Moving file system to the HSA <code>dm-switch</code>	94
Disabling the HSA	94
Chapter 9. Summary of changes applied on the host.	97
Changes applied on a Linux host	97
Multipath configuration	97
Multipath service(s)	98
iSCSI daemon configuration	98
iSCSI initiator name	98
iSCSI service	98
Changes applied on a Windows Server host	99
Removal of previous HAK versions	99
LUN0 device driver	99
Disk timeout setting	99
iSCSI service	99
IBM MPIIO Load Balancing service	99
IBM MPIIO Load Balancing agent	99

Built-in MPIO feature	100
Chapter 10. IBM storage systems and VMware ESXi hosts	101
Connectivity types for ESXi hosts.	101
Fibre Channel SAN zoning with ESXi hosts	101
iSCSI target discovery for ESXi hosts	102
Multipathing and load balancing with ESXi hosts	103
Notices	105
Trademarks	107
Index	109

Figures

1.	A host connected to the IBM storage system	2
2.	The workflow	3
3.	Preparation	5
4.	Fibre Channel host connectivity	6
5.	iSCSI host connectivity	8
6.	FC zoning: single initiator with multiple targets.	9
7.	Backview arrow in the XIV GUI.	11
8.	Viewing FC port details	12
9.	System mini-dashboard	13
10.	Viewing FC port parameters	14
11.	FC port parameters	15
12.	XIV GUI - Hosts and Clusters - iSCSI Connectivity	17
13.	XIV GUI - iSCSI connectivity table	17
14.	XIV GUI - Define IP Interface - iSCSI	17
15.	System mini-dashboard	19
16.	Displaying iSCSI port parameters	19
17.	iSCSI port parameters	20
18.	Software installation	25
19.	Installation language selection panel	29
20.	Installation wizard	29
21.	Installation Wizard – Ready to install the program	30
22.	Single host attachment	37
23.	XIV GUI - Hosts and Clusters	44
24.	XIV GUI - Hosts and Clusters table	44
25.	XIV GUI - Add host dialog box	45
26.	XIV GUI - Add Port dialog box	46
27.	New Host option	47
28.	Add Host form	48
29.	XIV GUI - Modifying LUN Mapping to a host	49
30.	XIV GUI - Modifying LUN Mapping	50
31.	Volume mini-dashboard	50
32.	Adding new volume mapping to host.	51
33.	Adding new volume mapping for a host.	52
34.	Cluster host attachment	53
35.	A cluster connected to the IBM storage systems	54
36.	Add Cluster dialog box	55
37.	XIV GUI - Cluster added to the Hosts and Clusters table	55
38.	XIV GUI - Add to Cluster.	56
39.	XIV GUI - Add Host(s) to Cluster dialog box	56
40.	XIV GUI - Host added to a cluster	57
41.	New Cluster option	58
42.	Create Cluster form	59
43.	Hosts & Clusters table	59
44.	XIV GUI - Modify LUN Mapping to a cluster	60
45.	XIV GUI - Modify LUN Mapping	61
46.	Right-clicking a specific host of a cluster	61
47.	Volume mini-dashboard	62
48.	Adding new volume mapping to cluster	63
49.	Adding new volume mapping for a cluster	64
50.	Additional tasks	65
51.	XIV GUI - Hosts Connectivity	90
52.	XIV GUI - Hosts Connectivity table	90
53.	Hosts & Clusters Views menu	91
54.	Hosts Connectivity table	92
55.	A host connected to IBM Spectrum Accelerate	93

56.	SAN Fibre Channel zoning layout	102
57.	Connecting ESXi servers over iSCSI	103

Tables

1.	Primary settings for proper iSCSI connectivity	16
2.	Optional CLI arguments for unattended installation	31
3.	xiv_attach exit codes	43
4.	Optional arguments for xiv_syslist	67
5.	xiv_syslist exit codes.	69
6.	Optional arguments for xiv_devlist	71
7.	xiv_devlist exit codes.	73
8.	Arguments for xiv_fc_admin	74
9.	xiv_fc_admin exit codes	75
10.	Arguments for xiv_iscsi_admin	76
11.	xiv_iscsi_admin exit codes	78
12.	xiv_diag exit codes	84
13.	Host profile information parameters	85
14.	Arguments for xiv_host_profiler	86
15.	xiv_host_profiler exit codes	88
16.	Initiators and targets in the SAN Fibre Channel zoning layout	102

About this guide

This guide describes how to prepare for, install, and use the IBM® Storage Host Attachment Kit (HAK) on supported operating systems.

Note: The last chapter provides reference information about how to physically connect VMware ESX or ESXi hosts to the IBM storage systems. However, the IBM Storage Host Attachment Kit software solution is not available for VMware ESX or ESXi hosts, and cannot be used with these hosts.

Who should use this guide

This guide is intended for system administrators who are familiar with the IBM storage systems and with the specific operating system that is used on their host(s).

Conventions used in this guide

The following notices are used to highlight key information in this guide.

Note: These notices provide important tips, guidance, or advice.

Important: These notices provide information or advice that might help you avoid inconvenient or difficult situations.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage can occur.

AIX

AIX users

Note: These notices provide information that is relevant only to the IBM AIX® operating system.

HP-UX

HP-UX users

Note: These notices provide information that is relevant only to the HP UNIX operating system.

Linux

Linux users

Note: These notices provide information that is relevant only to the Linux operating system.

Solaris

Solaris users

Note: These notices provide information that is relevant only to the Oracle Solaris operating system.

Note: These notices provide information that is relevant only to the Microsoft Windows Server operating system.

Related information and publications

You can find additional information and publications related to the IBM Storage Host Attachment Kit on the following information sources.

- IBM Knowledge Center(ibm.com/support/knowledgecenter)
- IBM XIV[®] Storage System on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJTAG)
- IBM Spectrum Accelerate on IBM Knowledge Center(ibm.com/support/knowledgecenter/STZSWD)
- IBM FlashSystem[®] A9000 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKMM)
- IBM FlashSystem A9000R on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKN5)
- *IBM XIV Storage System – Host Attachment and Interoperability*(www.redbooks.ibm.com/abstracts/sg247904.html)
- *IBM XIV Storage System – IBM Hyper-Scale Mobility Overview and Usage* (describing the XIV online volume migration feature) (www.redbooks.ibm.com/redbooks.nsf/redpieceabstracts/redp5007.html)

Depending on the operating system you are using on the host, refer to the following publications and information sources.

Operating system	Related publications and information sources
AIX	IBM Power Systems™ software website for AIX (www.ibm.com/systems/power/software/aix)
HP-UX	HP-UX 11i website (h71028.www7.hp.com/enterprise/us/en/os/hpux11i-overview.html)
Linux	<ul style="list-style-type: none"> • IBM and Linux website (www.ibm.com/linux) • Linux by RedHat (www.redhat.com/rhel) • Linux by SUSE (www.suse.com)
Solaris	<ul style="list-style-type: none"> • Oracle Solaris Patching Documentation Center (www.oracle.com/technetwork/systems/patches/solaris) • Oracle Solaris documentation (www.oracle.com/technetwork/documentation/index.html#sys_sw) • Oracle Solaris Online Forum (www.oracle.com/us/products/servers-storage/solaris)
Windows Server	<ul style="list-style-type: none"> • Microsoft TechNet website for Windows Server (technet.microsoft.com/en-us/windowsserver) • Microsoft Windows Server TechCenter (technet.microsoft.com/en-us/library/bb625087.aspx)

Getting information, help, and service

If you need help, service, technical assistance, or want more information about IBM products, you can find various sources to assist you. You can view the following websites to get information about IBM products and services and to find the latest technical information and support.

- IBM website (ibm.com[®])
- IBM Support Portal website (www.ibm.com/storage/support)
- IBM Directory of Worldwide Contacts website (www.ibm.com/planetwide)

IBM Publications Center

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The IBM Publications Center website (ibm.com/shop/publications/order) offers customized search functions to help you find the publications that you need. You can view or download publications at no charge.

Sending your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this guide or any other IBM Storage Host Software documentation:

- Go to the online feedback form (<http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strghosts.doc/icfeedback.htm>). You can use this form to enter and submit comments.
- You can send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
 - Exact publication title and version
 - Publication form number (for example: GA00-1111-22)
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Note: When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Chapter 1. Introduction

The IBM Storage Host Attachment Kit (HAK) is a software pack that simplifies the task of connecting a host computer to IBM XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000R and IBM FlashSystem A9000 storage systems.

The HAK provides a set of command-line interface (CLI) tools that help you perform different host-side tasks, such as: detect a physically connected IBM storage system (single system or multiple systems), detect volumes, define the host on the storage system, run diagnostics, and apply best practice native multipath connectivity configuration on the host.

After the host connection is established, storage volumes can be mapped to the host without any additional configuration, and can be accessed and used from the host for a range of storage operations.

Why use the host attachment kit?

Host I/O operations must be equitably distributed among the interface modules of a storage system.

Although the storage system distributes the traffic between its I/O and data modules, it is the responsibility of the storage administrator to ensure that host I/O operations are equitably distributed among the interface modules of the storage device. This workload balance must be monitored and assessed over time when host traffic patterns change.

Accordingly, each host must be properly configured before it can perform I/O operations on the storage system, and the HAK provides the host administrator with all the required tools for automatic and simpler host configuration. In addition, the HAK facilitates the monitoring and management of storage volumes from the host.

Connectivity types

You can connect hosts to the IBM storage systems over Fibre Channel (FC) or iSCSI connections, depending on the following conditions:

- The operating system used on the host. Some operating systems support only a specific connection type (for more information, refer to the HAK release notes issued for the operating system).
- The connectivity hardware that is installed or available on the host. FC connectivity requires a host bus adapter (HBA), and iSCSI connectivity requires an Ethernet or Gigabit Ethernet port.

You can connect hosts to Spectrum Accelerate systems only over iSCSI connections.

Concept diagram

Refer to the following graphic illustration to better grasp the concept of host connection.

Figure 1 on page 2 illustrates how a host on which the HAK is installed connects over Fibre Channel (FC) or iSCSI to the IBM storage system. The host may be

using IBM AIX, HP UNIX, Linux, Oracle Solaris, or Microsoft Windows Server as its operating system, which can either be installed locally or booted from a storage area network (SAN). The HAK helps detect the storage system, as well as any volume that is available for use by the host.




Figure 1. A host connected to the IBM storage system

Note: The HAK does not have to be installed locally on the host. For more information, refer to Chapter 3, “Software installation,” on page 25.

Compatibility and requirements

For the complete and up-to-date information about the compatibility and requirements of the IBM Storage Host Attachment Kit (HAK), refer to the latest release notes issued for the HAK version and the specific operating system that you are using.

You can obtain the latest release notes in the “Platform and application integration” section of the documentation for supported storage products, such as IBM XIV, IBM Spectrum Accelerate, IBM FlashSystem A9000 and IBM FlashSystem A9000R.

The workflow

The workflow includes preparation, installation, host attachment, and additional tasks.

Each stage in this workflow has its own chapter in this guide, as illustrated in Figure 2.




Figure 2. The workflow

Download site

The IBM Storage Host Attachment Kit is available for download from the IBM Fix Central(www.ibm.com/support/fixcentral).

Before you proceed

Before you proceed to the Preparation stage in the workflow, make sure that you have access to a workstation on which your preferred IBM storage system management software is locally installed or accessible via a remote connection.

You can manage the connected storage system using the GUI (graphical user interface) or CLI (command-line interface).

Note: Use the XIV GUI to manage the XIV and Spectrum Accelerate storage systems. However, for the FlashSystem A9000 and A9000R, the Hyper-Scale Management application is used.

For additional assistance and more information on how to obtain, install, and use the storage system management software, refer to the relevant section on IBM Knowledge Center, as listed in “Related information and publications” on page xii.

Chapter 2. Preparation

Before you can install and use the IBM Storage Host Attachment Kit (HAK), preparation actions might be required.




Figure 3. Preparation

The preparation stages are:

- Stage 1: Physically connecting the host to the IBM storage system
- Stage 2: Setting connectivity parameters
- Stage 3: Preparing the host

Note: Perform these preparation stages only once per host or cluster. If you have already performed these preparation stages, skip to Chapter 3, “Software installation,” on page 25.

Stage 1: Physically connecting the host to the IBM storage system

Any host on which you intend to use the HAK must be physically connected to the IBM storage system.

Depending on the hardware configuration and the operating system that is used on the host, use one of the following options to physically connect your host to the IBM storage system:

- Option 1: Using Fibre Channel (FC) connections
- Option 2: Using iSCSI connections

Attention: Although you can use both FC and iSCSI to connect different hosts to the same IBM storage system that supports both connection methods, use only one connection type per single host or cluster to avoid possible issues.

Option 1: Using Fibre Channel (FC) connections

The Fibre Channel (FC) protocol (FCP) is used to establish connectivity over FC infrastructure.

Hosts can be connected to the storage system via a private FC storage fabric, through an FC switch (not via a direct connection).

For the physical connection on the host side, install at least one host bus adapter (HBA) on the host. The HBA provides physical FC ports to which you can connect FC cables. The IBM Storage Host Attachment Kit supports different HBA brands and types. For the latest support information and compatibility matrix, refer to the IBM System Storage® Interoperation Center website (www.ibm.com/systems/support/storage/config/ssic).

Note:

- Fibre Channel is supported by all the operating systems for which the Host Attachment Kit is released (AIX, HP-UX, Linux, Solaris, or Windows Server).
 - For best performance, install the latest firmware and drivers for the HBAs that are in use. The HBA vendor should provide the latest firmware and drivers.
-

Figure 3 illustrates how a single host is connected to an IBM storage system using an FC switch and FC connections (FCP stands for Fibre Channel Protocol).

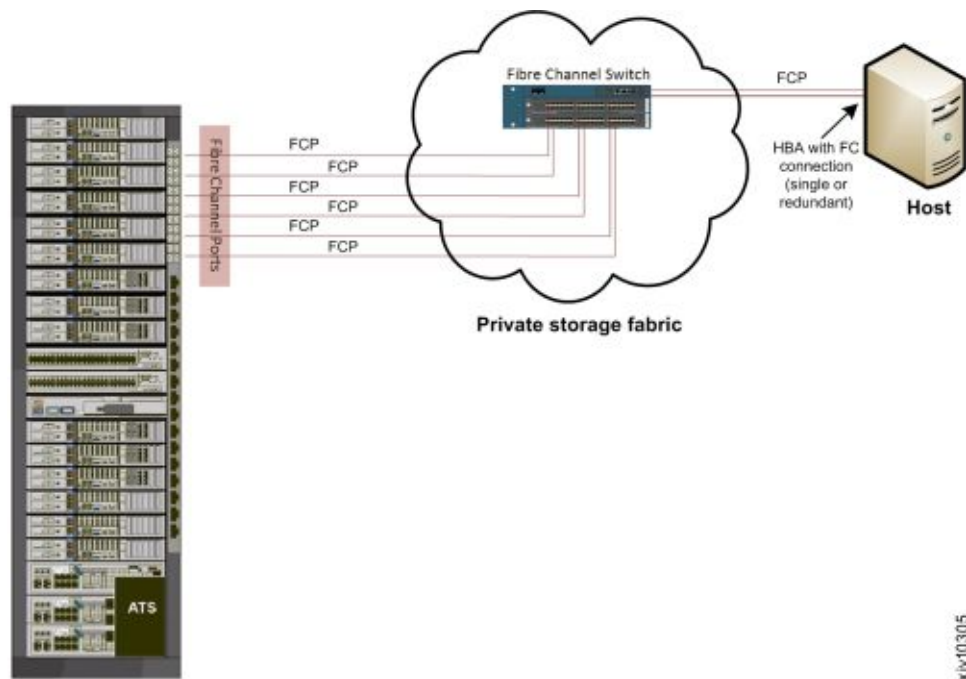


Figure 4. Fibre Channel host connectivity

Important: A high availability (HA) configuration is the best practice for FC connectivity. For best performance, use a dedicated FC connectivity infrastructure.

For additional information on how to physically connect your host(s) to the supported storage system over Fibre Channel, refer to the following documents:

- Network requirements chapter in *IBM XIV Planning Guide*, available in the IBM XIV Storage System documentation(www-01.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).
- Network requirements chapter in *IBM FlashSystem A9000 Planning Guide*, available in the IBM FlashSystem A9000 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKMM).
- Network requirements chapter in *IBM FlashSystem A9000R Planning Guide*, available in the IBM FlashSystem A9000R on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKN5).
- *IBM XIV Storage System – Host Attachment and Interoperability*(www.redbooks.ibm.com/abstracts/sg247904.html), available on the IBM Storage Redbooks® website (www.redbooks.ibm.com/redbooks.nsf/portals/storage).

Option 2: Using iSCSI connections

The iSCSI protocol is used to establish connectivity over Ethernet infrastructure.

Hosts can be connected to the storage system via a private Ethernet-based LAN or WAN, through an Ethernet switch (not via a direct connection).

AIX	<p>AIX users</p> <p>Note: The IBM Storage Host Attachment Kit for AIX does not support iSCSI connectivity. Only FC connectivity is supported on AIX (see Option 1: Using FC connections).</p>
HP-UX	<p>HP-UX users</p> <p>Note: The IBM Storage Host Attachment Kit for HP-UX does not support iSCSI connectivity. Only FC connectivity is supported on HP-UX (see Option 1: Using FC connections).</p>

Figure 4 illustrates how a single host is connected to the IBM storage system using an Ethernet switch and Gigabit Ethernet (GbE) connections.

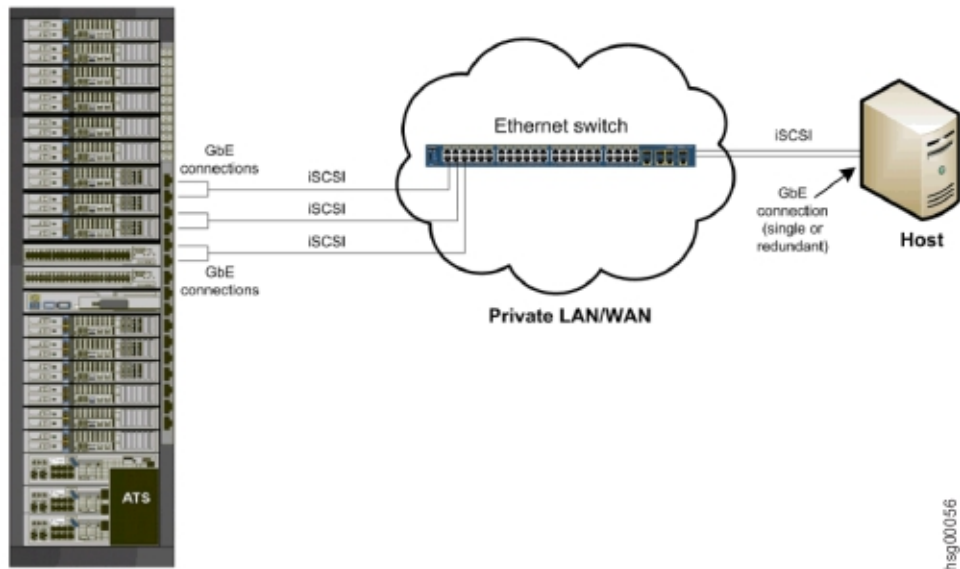


Figure 5. iSCSI host connectivity

Important: A high availability (HA) configuration is the best practice for iSCSI connectivity. For best performance, use a dedicated iSCSI connectivity infrastructure.

For additional information on how to physically connect your host(s) to the IBM storage system over Ethernet-based networks, refer to the following documents:

- Network requirements chapter in *IBM XIV Planning Guide*, available in the IBM XIV Storage System documentation (www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).
- Network requirements chapter in *IBM Spectrum Accelerate Planning Guide*, available in the IBM Spectrum Accelerate documentation (http://www.ibm.com/support/knowledgecenter/STZSWD/landing/IBM_Spectrum_Accelerate_welcome_page.html).
- Network requirements chapter in *IBM FlashSystem A9000 Planning Guide*, available in the IBM FlashSystem A9000 on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKMM).
- Network requirements chapter in *IBM FlashSystem A9000R Planning Guide*, available in the IBM FlashSystem A9000R on IBM Knowledge Center (ibm.com/support/knowledgecenter/STJKN5).
- *IBM XIV Storage System – Host Attachment and Interoperability* (www.redbooks.ibm.com/abstracts/sg247904.html), available on the IBM Storage Redbooks website (www.redbooks.ibm.com/redbooks.nsf/portals/storage).

Stage 2: Setting connectivity parameters

After you physically connect your host or cluster to the IBM storage system, you must ensure that the connection type (FC or iSCSI) is properly defined on that system.

Depending on the connection type that you are using, refer to the following subsections:

- Overview of Fibre Channel zoning
- Identifying and checking Fibre Channel ports from the XIV GUI
- Identifying and checking Fibre Channel ports from the IBM Hyper-Scale Manager
- Overview of iSCSI settings
- Configuring iSCSI ports from the XIV GUI
- Configuring iSCSI ports from the IBM Hyper-Scale Manager

Overview of Fibre Channel zoning

Zoning is a connection isolation technique that is required when connecting hosts with the IBM storage system over Fibre Channel SAN.

Zoning achieves the following:

- Creates a barrier between different network environments. Only the members of the same zone can communicate within that zone, and all external communications are blocked.
- Isolates any single host bus adapter (HBA) for security and reliability reasons.
- Allows finer segmentation of the switched fabric.

With proper zoning, you can avoid the following situations:

- HBAs of different vendors operate differently upon error recovery. This may impact other hosts connected to the same switch if those are not isolated through zoning.
- Any change in the SAN fabric, such as a change caused by a server restart or a new product being added to the SAN – triggers a registered state change notification (RSCN). An RSCN requires any change-aware device to acknowledge the change while interrupting its own traffic flow.

The most secure zoning is to have zones that consist of a single initiator and multiple targets, as illustrated in Figure 6 (see zones 1, 2, 3, and 4).




Figure 6. FC zoning: single initiator with multiple targets

Attention: Do not configure more than 24 logical paths per host, as this may impact overall stability. There is no advantage in configuring more than 24 logical paths.

For additional information about zoning, refer to *IBM XIV Storage System – Host Attachment and Interoperability* (www.redbooks.ibm.com/abstracts/sg247904.html), available on the IBM Storage Redbooks website (www.redbooks.ibm.com/redbooks.nsf/portals/storage).

Identifying and checking Fibre Channel ports on the IBM storage system

Before you can set the FC zoning for your hosts and IBM storage system (see Overview of Fibre Channel zoning), the FC ports on the storage system must be identified.

Fibre Channel ports are identified by their World Wide Port Name (WWPN).

Depending on the storage system type that you are using, refer to the following subsections:

- Identifying and checking Fibre Channel ports from the XIV GUI
- Identifying and checking Fibre Channel ports from the IBM Hyper-Scale Manager

Identifying and checking Fibre Channel ports from the XIV GUI

Before you can set the FC zoning for your hosts and storage system (see Overview of Fibre Channel zoning), the FC ports on the storage system must be identified.

Fibre Channel ports are identified by their World Wide Port Name (WWPN).

Use the XIV GUI or the XCLI (XIV CLI) to determine the FC port WWPNs, and also to check that the FC ports are set up correctly.

To identify the FC ports from the XIV GUI:

1. Log in to the IBM XIV GUI management application and zoom into the storage system that you want to configure.
2. From the main view of an IBM XIV Storage System, click the **Backview** arrow at the bottom of the XIV storage system to reveal the patch panel on the back.



Figure 7. Backview arrow in the XIV GUI

The back patch panel is displayed.

3. Move the mouse pointer over any particular FC port to reveal the port details, including the WWPN.



Figure 8. Viewing FC port details

Alternatively, to identify the FC ports from the XIV CLI (XCLI), enter the following command:

```
fc_port_list [ module=ComponentId | fcport=ComponentId ]
```

Note:

- The WWPNs of an XIV storage system are static. The last two digits of the WWPN indicate from which module and port the WWPN came.
- By default, port number 4 of each interface module is configured as an initiator.

Identifying and checking Fibre Channel ports from the IBM Hyper-Scale Manager

Before you can set the FC zoning for your hosts and storage system (see Overview of Fibre Channel zoning), the FC ports on the storage system must be identified.

Fibre Channel ports are identified by their World Wide Port Name (WWPN).

Use the IBM Hyper-Scale Manager to determine the FC port WWPNs, and also to check that the FC ports are set up correctly.

To identify the FC ports from the IBM Hyper-Scale Manager application:

1. Log in to the IBM Hyper-Scale Manager application and select a storage system that you want to configure.

- From the **System** mini-dashboard, click the **P.Panel** relation icon to display the list of FC ports on the system patch panel.

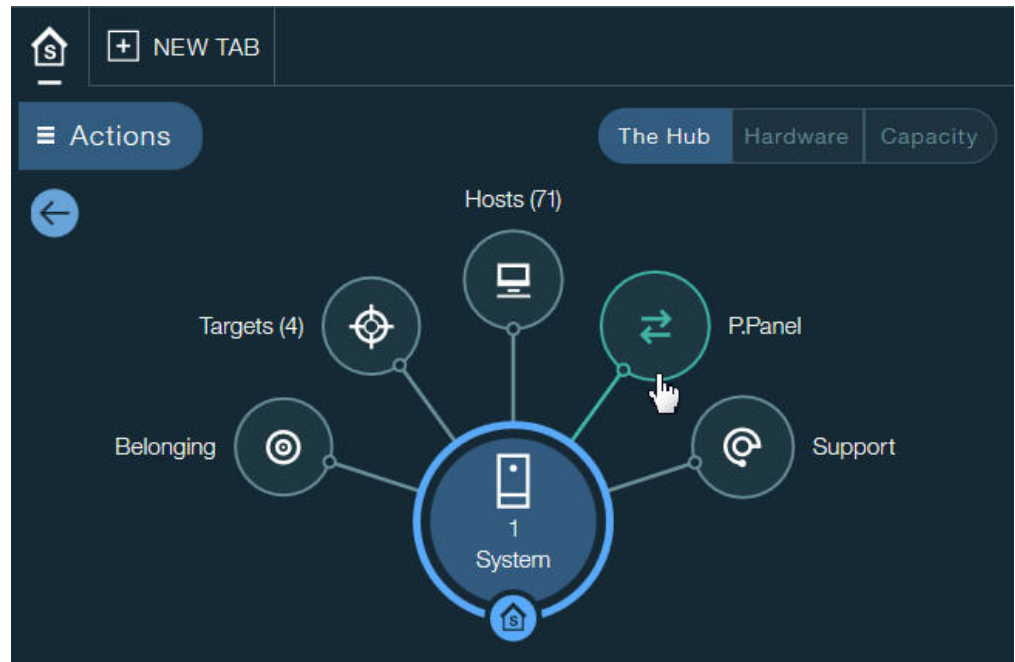


Figure 9. System mini-dashboard

The list of FC ports is displayed.

- Click the **Actions** button on any particular FC port and select **View/Edit FC Port** to reveal the port parameters, including the WWPN.

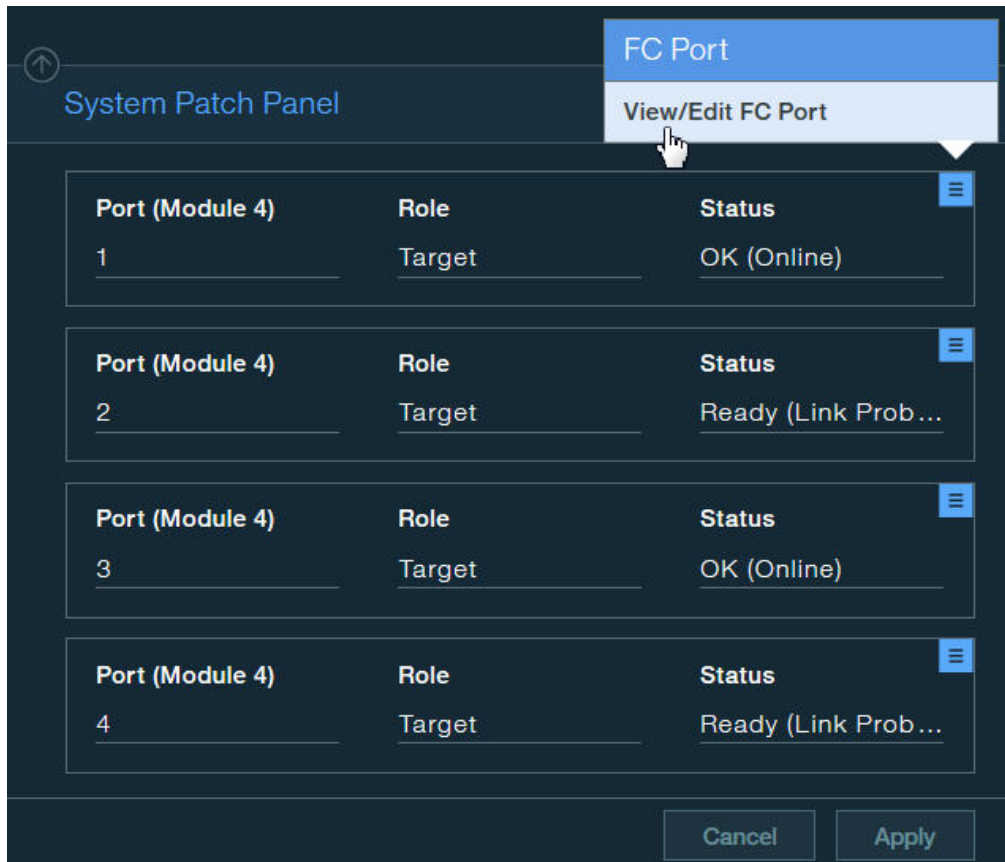


Figure 10. Viewing FC port parameters

The FC port parameters are displayed.

4. In the **FC Port** form, identify the WWPN of the port.

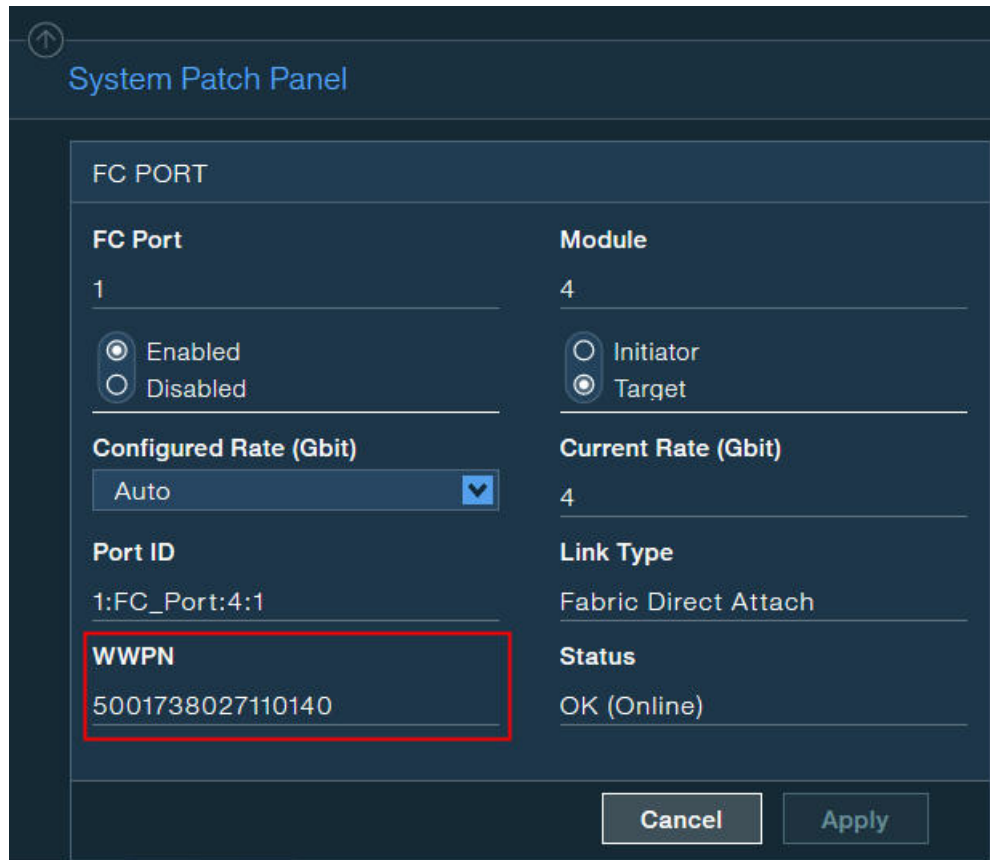


Figure 11. FC port parameters

Overview of iSCSI settings

The IBM storage system acts as a Transmission Control Protocol (TCP) server for iSCSI connections.

All iSCSI packets are routed through the Ethernet port from which the iSCSI connection was initiated.

AIX	<p>AIX users</p> <p>Note: The IBM Storage Host Attachment Kit for AIX does not support iSCSI connectivity. Only FC connectivity is supported on AIX (see Overview of Fibre Channel zoning).</p>
HP-UX	<p>HP-UX users</p> <p>Note: The IBM Storage Host Attachment Kit for HP-UX does not support iSCSI connectivity. Only FC connectivity is supported on HP-UX (see Overview of Fibre Channel zoning).</p>

The following table specifies the primary settings that must be configured on the storage system for proper iSCSI connectivity with the host.

Table 1. Primary settings for proper iSCSI connectivity

Setting type	Configuration on the storage system
iSCSI qualified name (IQN)	<p>Required for each host that uses iSCSI to connect to the storage system. The IQN uniquely identifies the different entities that utilize iSCSI on the IP network, and enables the operation of the iSCSI initiator.</p> <p>The IQN of the storage system is a pre-configured factory setting that cannot be changed.</p>
IP interface	<p>iSCSI uses IP-based communication. On the storage system, each iSCSI port is defined as an IP interface with its own IP address. Alternatively, the ports can be bundled (bonded) for load balancing to a single logical iSCSI port with a single IP address.</p> <p>By default, six iSCSI target ports are predefined on the storage system to serve hosts over iSCSI.</p>
Default gateway IP address	<p>Required if the connected hosts do not reside on the same Layer-2 subnet.</p>
Maximum transmission unit (MTU)	<p>Required if your network supports an MTU that is larger than the standard one. The MTU default value is 1,536 bytes and the maximum value is 8,192 bytes.</p> <p>You should specify the largest possible MTU. If supported by the switches and routers, use up to 8,192 bytes.</p>

Depending on the storage system type that you are using, refer to the following subsections:

- Configuring iSCSI ports from the XIV GUI
- Configuring iSCSI ports from the IBM Hyper-Scale Manager

Configuring iSCSI ports from the XIV GUI

Perform the following procedure to configure iSCSI ports on the storage system from the XIV GUI.

Procedure

1. Log in to the IBM XIV GUI management application and zoom into the storage system that you want to configure.
2. Move the mouse over the **Hosts and Clusters** icon and then click **iSCSI Connectivity**.




Figure 12. XIV GUI - Hosts and Clusters - iSCSI Connectivity

The iSCSI Connectivity table is displayed.

Name	Address	Netmask	Gateway	MTU	Module	Ports
iscsi_1_9_4	172.16.2.31	255.248.0.0	0.0.0.0	1500	1:Module:3	4
iscsi_1_9_3	172.16.2.30	255.248.0.0	0.0.0.0	1500	1:Module:3	3
iscsi_1_9_2	172.16.2.29	255.248.0.0	0.0.0.0	1500	1:Module:3	2
iscsi_1_9_1	172.16.2.28	255.248.0.0	0.0.0.0	1500	1:Module:3	1
iscsi_1_8_4	172.16.2.27	255.248.0.0	0.0.0.0	1500	1:Module:3	4
iscsi_1_8_3	172.16.2.26	255.248.0.0	0.0.0.0	1500	1:Module:3	3
iscsi_1_8_2	172.16.2.25	255.248.0.0	0.0.0.0	1500	1:Module:3	2

Figure 13. XIV GUI - iSCSI connectivity table

- Click the **Define** button located at the top of the window. The Define IP Interface – iSCSI dialog box is displayed.

Figure 14. XIV GUI - Define IP Interface - iSCSI

- Fill in the following fields (for more information, see Table 1 on page 16):

- **Name** – Name of the iSCSI interface.
- **Address** – IP address of the iSCSI interface.
- **Netmask** – Enter the netmask address, or select an existing one from the list box.
- **Default Gateway** – If needed, enter the IP address of the default gateway.
- **MTU** – Maximum transmission unit. The default is 4500; change this value only if needed.
- **Module** – Select the physical module on which the iSCSI interface should be set.
- **Port Number** – Select the physical port number on the previously selected physical module.

Note: All devices in a network must use the same MTU. If in doubt, set MTU to 1500, because 1500 is the default value for Gigabit Ethernet. Performance might be impacted if the MTU is set incorrectly.

5. Click **Define**. The iSCSI interface is set. Alternatively, you can use the following XCLI command to create a new IP interface for iSCSI:

```
ipinterface_create ipinterface=IPInterfaceName
address=IPaddress netmask=NetworkMask
[ gateway=DefaultGateway ] [ mtu=MTU ] module=ComponentId ports=P1[,P2]
```

For more information about the **ipinterface** command, refer to the IBM XIV XCLI Reference Guide, available in the IBM XIV Storage System documentation (www-01.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).

Configuring iSCSI ports from the IBM Hyper-Scale Manager

Perform the following procedure to configure iSCSI ports from the IBM Hyper-Scale Manager application.

Procedure

1. Log in to the IBM Hyper-Scale Manager application and select a storage system that you want to configure.
2. From the **System** mini-dashboard, click the **P.Panel** relation icon to display the list of iSCSI ports on the system patch panel.

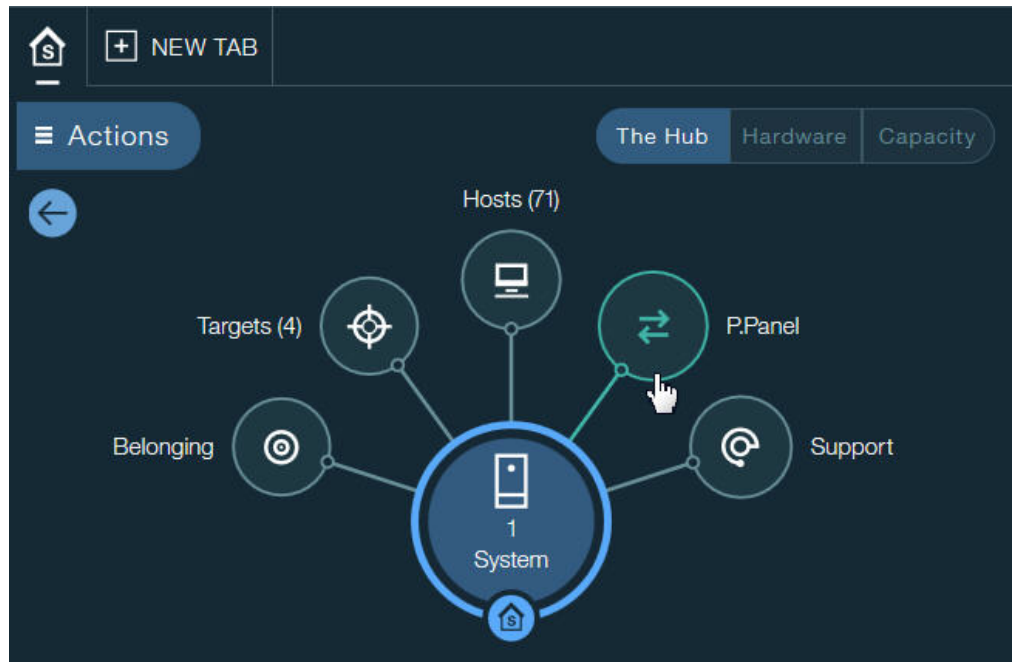


Figure 15. System mini-dashboard

The list of iSCSI ports is displayed.

3. Click the **Actions** button on any particular iSCSI port and select **Define IP Interface** to display the port parameters.


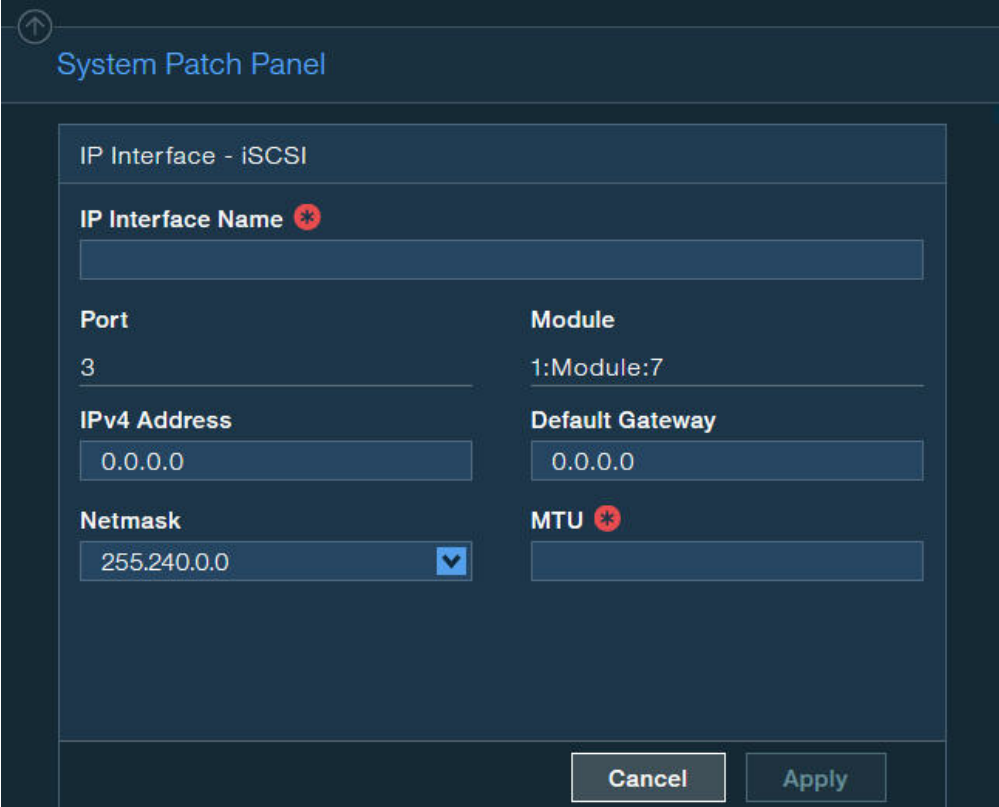


Figure 16. Displaying iSCSI port parameters

The iSCSI port parameters are displayed.

4. In the **IP Interface - iSCSI** form, fill in the following fields (for more information, see Table 1 on page 16):
 - **IP Interface Name** – Define the name of the iSCSI interface.
 - **Port** – Select the physical port number on the previously selected physical module.
 - **Module** – Select the physical module on which the iSCSI interface should be set.
 - **IPv4 Address** – IP address of the iSCSI interface.
 - **Default Gateway** – If needed, enter the IP address of the default gateway.
 - **Netmask** – Enter the netmask address, or select an existing one from the list box.
 - **MTU** – Select the maximum transmission unit value. The default is 4500; change this value only if needed.



The screenshot shows a 'System Patch Panel' window with a dark blue header. Below the header is a form titled 'IP Interface - iSCSI'. The form contains several input fields and dropdown menus. The 'IP Interface Name' field is empty and has a red asterisk icon to its right. The 'Port' field contains the value '3'. The 'Module' field contains the value '1:Module:7'. The 'IPv4 Address' field contains '0.0.0.0'. The 'Default Gateway' field contains '0.0.0.0'. The 'Netmask' field contains '255.240.0.0' and has a dropdown arrow on its right. The 'MTU' field is empty and has a red asterisk icon to its right. At the bottom right of the form are two buttons: 'Cancel' and 'Apply'.

Figure 17. iSCSI port parameters

Note: All devices in a network must use the same MTU. If in doubt, set MTU to 1500, because 1500 is the default value for Gigabit Ethernet. Performance might be impacted if the MTU is set incorrectly.

5. Click **Apply**. The iSCSI interface is set.

Using CHAP authentication for iSCSI connectivity

The IBM storage systems support industry-standard unidirectional iSCSI CHAP authentication, allowing identity validation of the iSCSI initiator that attempts to log in to the storage system.

Important: CHAP authentication pre-configuration on the storage system is defined per individual host. There is no global CHAP configuration that can apply to multiple hosts. By default, hosts are defined without CHAP authentication.

To set up CHAP authentication for a host, you must configure the following two parameters on the storage system:

- CHAP name – The username that the host must use for CHAP authentication.
- CHAP secret – The password that the host must use for CHAP authentication.

CHAP name and secret parameter guidelines

The following guidelines apply to the CHAP name and secret parameters:

- The CHAP secret must contain 96 bits to 128 bits (12 to 16 ASCII characters).
- Both the CHAP name and CHAP secret must be set together if CHAP authentication is to be used. Authentication will not work if only one parameter (name or secret) is defined.

CLI commands for configuring CHAP authentication

Use the following CLI commands to configure CHAP authentication.

To define a new host with CHAP parameters:

```
host_define host=[hostName] iscsi_chap_name=[chapName]
iscsi_chap_secret=[chapSecret]
```

Alternatively, if the host already exists:

```
host_update host=[hostName] iscsi_chap_name=[chapName]
iscsi_chap_secret=[chapSecret]
```

To clear the CHAP parameters (if no longer needed):

```
host_update host=[hostName] iscsi_chap_name= iscsi_chap_secret=
```

Note:

- For information about how to define the CHAP authentication using the GUI, see “Defining a host from the XIV GUI” on page 44.
 - For information about how to use the CHAP name and secret when creating a new iSCSI target on the host, see “Setting iSCSI connectivity parameters” on page 75.
-

Stage 3: Preparing the host

Prepare your host for attachment as necessary.

The host preparation stage includes the following tasks:

- Validating the system requirements
- Validating your admin user privileges
- Optional: booting the operating system from a storage area network (SAN)
- Preparing AIX hosts (if applicable)

Validating the system requirements

Each operating system requires different validations.

Before installing the IBM Storage Host Attachment Kit (HAK) on the host, verify that the installed operating system meets the requirements specified in the *Compatibility and requirements* section of the latest release notes, issued for the HAK version that you are using.

Windows Server	Windows Server users Important: <ul style="list-style-type: none">• Refer to the latest release notes for a list of Microsoft Hotfixes that must be installed manually prior to installing the HAK.• Ensure that Microsoft iSCSI Initiator (version 2.08 or later) is installed with the following components:<ul style="list-style-type: none">– Initiator software and service – included in Windows Server 2008 or later.– Microsoft MPIO multipathing support for iSCSI – in Windows Server 2008 or later, MPIO should be installed from the Server Manager.
Linux	Linux users Important: Refer to the latest release notes for a list of the iSCSI or Fibre Channel components that must be installed on your Linux operating system.
AIX	AIX users Important: <ul style="list-style-type: none">• The <code>bos.adt.libm</code> package must be installed on the host. For more information, refer to the IBM Fileset information website (www.ibm.com/support/docview.wss?uid=isg1fileset-1318926131).• To avoid possible host reboots and downtime when using the HAK, it is recommended to always use the latest AIX TL update on the host. These AIX updates have native support for the IBM storage system. For more information, refer to the release notes.

Validating your admin user privileges

You must perform the HAK installation and all other kit operations with **root** or **administrator** user privileges on the host.

Before installing and using the HAK, make sure that your user name and password enable you to log in with all the required privileges.

Windows Server

Windows Server users

Important: If User Account Control is enabled on Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012, you must run all CLI commands as an administrator. To bypass the account control limitations, go to the `C:\WINDOWS\System32` directory, right-click **cmd.exe**, and then select **Run as Administrator**.

Optional: booting the operating system from a storage area network (SAN)

The host operating system (OS) can be booted from a local disk installation or from a remote storage area network (SAN).

To learn how to boot the OS from a SAN, refer to *IBM XIV Storage System – Host Attachment and Interoperability* (www.redbooks.ibm.com/abstracts/sg247904.html), available on the IBM Storage Redbooks website (www.redbooks.ibm.com/redbooks.nsf/portals/storage).

Preparing AIX hosts (if applicable)

Refer to this section if you intend to use the built-in multipathing application (MPIO) of AIX 6.1.

AIX

Note: This section applies only to AIX 6.1.

Skip this section if you are using AIX 7.1 or a different operating system on the host.

The AIX versions detailed in the note above require pre-configuration before you can use the **xiv_attach** utility (for more information about **xiv_attach**, see Chapter 4, “Single host attachment,” on page 37).

Note: To avoid the need for pre-configuration, use the latest TL and service pack on your AIX host. The latest AIX updates have native support for the IBM storage system. For more information, refer to the release notes.

To determine if your AIX version requires pre-configuration, run the following command:

```
manage_disk_drivers -l
```

If the command output includes **2810XIV**, your AIX host can identify the storage system automatically. If the command returns empty output or output that does not include 2810, pre-configuration is required for your AIX version.

Refer to the following subsections for pre-configuring:

- “Queue depth”
- “Load balancing for Round-robin”
- “Multipathing”

Queue depth

The queue depth default value for TL10 and TL03 is 32; the default value for previous AIX versions is 1.

```
export XIV_DEF_QD=32
```

Load balancing for Round-robin

Load balancing for Round-robin applies only for AIX managed multipathing.

By default, the package configures multipathing to work in Round-robin mode for TL10 and TL03 and failover mode for previous versions. You can change the multipathing mode to failover-only:

```
export XIV_DEF_ALGO=fail_over
```

If you define `failover_only`, only one path is used for any given hard disk. For Tls earlier than 10 and 03 respectively, failover is the only method by which you can use queue depths higher than 1.

Multipathing

If you are using a version of AIX that does not automatically identify the IBM storage system and you want to use third-party multipathing software to manage the system, you must set the following environment variable **before** you run the **xiv_attach** utility:

```
export XIV_NOMP=yes
```

In versions later than AIX 6.1 TL03, you must run the following command to use a third-party multipathing solution:

```
manage_disk_drivers -d 2810XIV -o AIX_non_MPIO
```

Note:

- The third party multipathing application must be supported by the storage system.
 - The disks may be detected as '**IBM 2810XIV Non-MPIO Fibre Channel Disk**' and may not be compatible with multipathing as a result.
 - If a real volume is not assigned to LUN 0, it is not multipathed and it may be detected as '**IBM 2810XIV-LUN-0 Fibre Channel Array Controller**'.
-

Chapter 3. Software installation

After the preparation stages are complete, you can begin the software installation.




Figure 18. Software installation

The software installation package includes the following components:

- IBM Storage Host Attachment Kit (HAK) – includes the following command-line interface (CLI) utilities:
 - **xiv_attach**
 - **xiv_devlist**
 - **xiv_syslist**
 - **xiv_fc_admin** and **xiv_iscsi_admin** (**xiv_iscsi_admin** is not supported on AIX and HP-UX)
 - **xiv_detach** (applicable to Windows Server only)
 - **xiv_diag**
 - **xiv_host_profiler**
- Operating system hotfixes (applicable to Windows Server only)

The following options are available for the HAK software installation:

- Option 1: Performing a clean local installation
- Option 2: Upgrading from a previous local installation
- Option 3: Using the portable package without installing

For information about how to uninstall the HAK, see “Removing the Host Attachment Kit software” on page 33.

Option 1: Performing a clean local installation

You can perform a clean local installation on any new host that has no previous HAK installations, or on any existing host from which the previous HAK installation has been completely removed.

Note: On new hosts, it is recommended to always install the latest available HAK package, rather than installing an old version and then upgrading to the newer HAK package.

Important: Refer to the latest release notes for up-to-date information regarding compatibility, limitations, and requirements.

Depending on the operating system used on the host, refer to one of the following sections:

- “Installing on AIX, HP-UX, Linux, or Solaris”
- “Installing on Windows Server” on page 28

Installing on AIX, HP-UX, Linux, or Solaris

This section describes the procedure to install the HAK on AIX, HP-UX, Linux, or Solaris.

Procedure

Perform the following procedure to install the HAK on AIX, HP-UX, Linux, or Solaris, using the native command prompt interface.

AIX

AIX users

Attention:

- Running the installation package removes any customized settings for AIX hard disk devices that correspond to storage system disks.
- For the virtual I/O server configuration, the installation uses the **oem_setup_env** environment.

-
1. Copy the HAK installation package file to a local folder on the host.
 2. Extract the installation package file (* stands for the OS and architecture package):

```
# gunzip -c IBM_Storage_Host_Attachment_Kit_2.6.0*.tar.gz | tar xvf -
```

3. Go to the extracted directory (referred to as <HAK extracted directory> below) and then run **install.sh** to start the installer:

```
# cd <HAK extracted directory>
# ./install.sh
```

The installer welcome message is displayed:


```

Welcome to the IBM Storage Host Attachment Kit installer.
Required OS packages:
+-----+-----+
| RHEL/CentOS | SLES |
+-----+-----+
| device-mapper-multipath | multipath-tools |
| sg3_utils | sg3_utils |
+-----+-----+
Attention: The installation may FAIL if any required OS package is missing.

Optional OS packages (iSCSI support):
+-----+-----+
| RHEL/CentOS | SLES |
+-----+-----+
| iscsi-initiator-utils | open-iscsi |
+-----+-----+
Would you like to proceed and install the Host Attachment Kit? [y/n]:

```

4. Enter yes (or y) to begin the installation. The following is an example of the displayed output:

```

Please wait while the installer validates your existing configuration...
-----
Please wait, the Host Attachment Kit is being installed...
-----

```

Note: The installation directory is: `/opt/xiv/host_attach`

When the installation is complete, the following prompt is displayed:

```

Installation successful.

Please refer to the user guide for information about how to configure this host.

-----
The IBM Storage Host Attachment Kit includes the following utilities:

xiv_attach: Interactive wizard that configures the host and verifies its configuration
for connectivity with the IBM storage system.

xiv_devlist: Lists of all storage volumes that are mapped to the host, with general info
about non-XIV volumes.

xiv_syslist: Lists all IBM storage systems that are detected by the host.

xiv_diag: Performs complete diagnostics of the host and its connectivity with the
IBM storage system, and saves the information to a file.

xiv_fc_admin: Allows you to perform different administrative operations for
FC-connected hosts and IBM storage systems.

xiv_iscsi_admin: Allows you to perform different administrative operations for
iSCSI-connected hosts and IBM storage systems.

xiv_host_profiler: Collects host configuration information and performs a comprehensive
analysis of the collected information.
-----

```

AIX**AIX users**

Note: Following the package installation, storage disks are multipathed and detected by the `lsdev` command as: **IBM 2810XIV Fibre Channel Disks**. If a real volume is not assigned to LUN0, it is not multipathed and it may be detected as: **IBM 2810XIV-LUN-0 Fibre Channel Array Controller**.

Installing on Windows Server

This section describes the procedure of installing the IBM Storage Host Attachment Kit (HAK) on Microsoft Windows Server.

The HAK installation package for Windows Server provides an installation wizard that helps you install all the required files on the host, including some required system updates (for more information, refer to the latest release notes).

Windows Server**Windows Server users**

Note: When installing the HAK in cluster environments: To install the HAK in Windows Failover Clustering 2008 environments, you must first install the HAK on the passive nodes. After a successful installation on these nodes, you can perform failover of the resources to the nodes on which the HAK is already installed. You can then install the HAK on the remaining nodes.

Windows Server**Windows Server users**

Note: When installing the HAK on hosts with existing multipath solutions: The Veritas Storage Foundation Dynamic Multipathing (DMP) does not rely on the native multipath I/O (MPIO) capabilities of the Windows Server operating system. Instead, it provides its own custom multipath I/O solution. Because these two solutions cannot co-exist on the same host, perform the following procedure if you intend to use the Veritas solution:

1. Install the Veritas Storage Foundation package.
2. Restart the host.
3. Install the IBM Storage Host Attachment Kit.

If the version of the Veritas Storage Foundation that you are using does not support the IBM storage system, install the Device Driver Installation Package 1 (DDI-1) for Veritas Storage Foundation 5.1. This update supports the IBM storage systems.

Perform the following procedure to install the HAK from the Windows GUI.

1. Copy the HAK installation package file (**IBM_Storage_Host_Attachment_Kit_2.6.0*_for_Windows*.zip**) that matches your Windows Server architecture (x86 or x64) to a local folder.

2. Extract the files from the zip file and then run the setup executable. The installation language selection panel is displayed.




Figure 19. Installation language selection panel

3. Select your preferred language for the installation and then click **OK**. The installer files are initialized and then IBM Storage Host Attachment Kit Installation Wizard is displayed.



Figure 20. Installation wizard

4. Click **Next**. The Ready to Install the Program panel is displayed.




Figure 21. Installation Wizard – Ready to install the program

5. Click **Install**. The installation starts.

Note: The installation directory is: C:\Program Files\XIV\host_attach

6. When the installation completes, click **Finish**. The IBM Storage Host Attachment Kit is added to the list of installed Windows programs.

Installing from the Windows command prompt interface

If you prefer using CLI for the installation, you can use the Windows command prompt program (**cmd.exe**) to run the HAK setup executable. In addition, you can use this method to install the HAK with verbose logging and in different levels of unattended installation modes.

If you want to run the CLI setup with verbose logging (creates a local log file), run the installation executable in its folder location as follows:

```
<IBM_Storage_Host_Attachment_Kit_2.6.0*_for_Windows*.exe> v"/1*vx setup.log"
```

Note: The **/1*vx** argument is part of the Microsoft Windows Installer (**msiexec.exe**) syntax for verbose output. For more information, refer to the Microsoft MSDN Library website(msdn.microsoft.com/en-US/library/ee251019%28v=BTS.10%29.aspx).

Installing in unattended mode

If you want to run the installation in unattended mode to eliminate user interaction during the installation, run the installation executable in its folder location as follows:

```
<IBM_Storage_Host_Attachment_Kit_2.6.0*_for_Windows*.exe> /s /v"/qn"
```

If needed, you can replace or add arguments as detailed in Table 2.

Table 2. Optional CLI arguments for unattended installation

Argument	Use after <XIV_host_attachment_windows*.exe> to:
/q	Install in silent mode.
/qn	See no user prompt during the installation.
/qb	See only basic user prompt during the installation.
/qr	See reduced user prompt during the installation.
/qf	See the entire user prompt during the installation. This is the default option.
/norestart	Prevent any host restart after the installation.
/promptrestart	Prompt before any host restart during or after the installation.
/forcerestart	Force a restart of the host after the installation.

Option 2: Upgrading from a previous local installation

When a new HAK version becomes available, the already installed HAK version can be upgraded.

If a previous version of the HAK is locally installed on the host, extract the new installation package and run the new installation file (**install.sh** on AIX, HP-UX, Linux, and Solaris; **setup** executable on Windows) to start the upgrade procedure. The installer automatically updates and replaces all the existing HAK files on the host.

Note:

- Starting from version 2.6.0, IBM XIV Host Attachment Kit is renamed to IBM Storage Host Attachment Kit.
- Unless specified otherwise in the release notes, you do not have to uninstall the previous version, and there is no need to detach the host from the storage system prior to the upgrade.

Important: You must run the **xiv_attach** command after any upgrade (for more information, see “Stage 1: Using the CLI attachment utility (xiv_attach)” on page 37).

If you are upgrading in a command prompt (CLI) environment, the following message is displayed:

```
Please wait while the installer validates your existing configuration...
-----
NOTE: Another version of Host Attachment Kit is already installed.
The installer will replace the currently installed version.
Would you like to replace the installed Host Attachment Kit? [y/n]:
```

Attention: If you are upgrading a Microsoft Cluster Server (MSCS) cluster, move all resources to unused nodes before the upgrade, so that the passive nodes are upgraded first.

Option 3: Using the portable package without installing

The portable package allows you to use the HAK "out of the box", without having to install it locally on the host.

You can run all HAK utilities from a shared network drive or from a portable USB flash drive onto which the portable HAK package has already been extracted or installed.

Note: Any locally installed HAK package becomes a portable HAK package when it is copied or moved from its original directory (`host_attach`) to a different directory.

Important:

- Before using the portable package, make sure that all the latest operating system (OS) updates or hotfixes have been installed. You can use the `xiv_fc_admin -V` or `xiv_iscsi_admin -V` commands to verify this and be notified if any OS component is missing. The `xiv_attach`, `xiv_fc_admin -C`, and `xiv_iscsi_admin -C` commands automatically install missing updates, but you may have to install additional updates manually.
- For more information about these commands, see Chapter 6, "Additional tasks," on page 65.
- For more information about the required OS updates or hot fixes, refer to the latest release notes.

Refer to one of the following subsections for more information:

- "Using a mounted network drive"
- "Using a portable USB flash drive" on page 33

Using a mounted network drive

This section describes how to use the HAK from a mounted network drive.

Procedure

To use the portable HAK package from a network drive:

1. Extract the files from `<HAK_build_name>_Portable.tar.gz` (for AIX, HP-UX, Linux, Solaris) or `<HAK_build_name>_Portable.zip` (for Windows) into a shared folder on a network drive.
2. Mount the shared folder to each host computer on which you intend to use the HAK, so that the folder would be recognized and accessible as a network drive. You can now use the IBM Storage Host Attachment Kit on any host to which the network drive is mounted.

Note: Whenever a newer HAK version is installed on the network drive, all hosts to which that network drive was mounted have immediate and direct access to the new HAK version.

Using a portable USB flash drive

This section describes how to use the portable HAK package from a USB flash drive.

Procedure

To use the portable HAK package from a USB flash drive ("Disk On Key"):

1. Extract the files from **<HAK_build_name>_Portable.tar.gz** (for AIX, HP-UX, Linux, Solaris) or **<HAK_build_name>_Portable.zip** (for Windows) into a folder on the USB flash drive.
2. Plug the USB flash drive into any host on which you want to use the IBM Storage Host Attachment Kit, and then run any HAK utility from the plugged-in drive.

Removing the Host Attachment Kit software

Use the appropriate software removal method when needed.

If the HAK is locally installed on the host (as opposed to a portable HAK package), you can uninstall the kit in one of the following methods:

- Detaching the host and uninstalling the software – applicable to Windows Server only.
- Uninstalling the software without detaching the host

Attention: Portable HAK packages do not require the uninstallation procedure. You can delete the portable HAK directory (on the network drive or the USB flash drive) to uninstall it. For more information about the portable HAK, see "Option 3: Using the portable package without installing" on page 32.

Detaching the host and uninstalling the software

This section describes when and how to use the **xiv_detach** utility on Windows Server.

About this task

Windows Server

Note: This section is relevant to Windows Server users only. If you are using a different operating system, see "Uninstalling the software without detaching the host" on page 35.

Prior to the removal of the HAK software, you can use the **xiv_detach** utility to remove all multipath settings from the host.

Note:

- Uninstalling the HAK without using the detachment utility keeps all existing multipath settings on the host. If you choose not to use the detachment utility, the regular uninstallation (see step 5) only removes the HAK utilities.
 - The detachment utility does not affect the host definitions or volume mappings on the storage system side. The host and any volume mapping to the host will remain defined on the storage system side until those definitions are manually removed using the GUI or CLI management tools.
-

Important: Always perform the **xiv_detach** operation and uninstall the HAK before upgrading the operating system (not before upgrading the HAK). After the operating system has been upgraded, you can reinstall the HAK and reattach the host.

Procedure

Perform the following procedure to remove all multipath settings and uninstall the HAK from the host.

1. Run the **xiv_detach** command and enter **yes** (or **y**) when prompted to confirm the detachment operation, as shown in the following example:

```
> xiv_detach

Welcome to the IBM host detachment wizard, version 2.6.0.
This wizard will help you detach this host from the storage system.

The detachment will unconfigure multipathing connectivity settings for this host.
This host will not be able work with IBM storage systems until it is reattached
Are you sure you want to continue? [default: no ]: y
```

The connectivity type choice message is displayed.

```
Please specify the connectivity type: [f]c / [i]scsi : f
```

2. Specify the connectivity type for which you want to remove all existing multipath settings: enter **f** for Fibre Channel or **i** for iSCSI. The detachment confirmation message is displayed.

```
Attention: If you have storage volumes that are mapped to this host,
it is recommended to unmap these volumes before unconfiguring the host.
Do you want to continue with the host detachment? [default: no ]: y
```

3. Enter **yes** (or **y**) to start the detachment.

```
Please wait while the host configuration is removed...
The host is now unconfigured for connectivity to the storage system.

The host has been successfully detached.

Press [ENTER] to exit.
```

4. Press Enter.
5. Go to **Control Panel > Programs > Programs and Features**, select the IBM Storage Host Attachment Kit from the list of installed programs, and then click **Uninstall/Change**. The uninstallation wizard guides you through the remaining uninstallation steps.

Note: The uninstall command removes all the installed HAK utilities, but it does not remove operating system hotfixes that were installed during the host attachment procedure.

Uninstalling the software without detaching the host

The regular uninstallation removes the locally installed HAK software without detaching the host, keeping all multipathing connections to the IBM storage system.

Attention: Portable HAK packages do not require the uninstallation procedure. You can delete the portable HAK directory (on the network drive or the USB flash drive) to uninstall it. For more information about the portable HAK, see “Option 3: Using the portable package without installing” on page 32.

To uninstall the HAK software from an AIX, HP-UX, Linux, or Solaris host, run the following command:

```
# /opt/xiv/host_attach/bin/uninstall
```

Displaying CLI manual pages

On supported operating systems that have a Unix-like command prompt environment, you can display help for the different HAK utilities by using the standard **man** command prompt option.

AIX

HP-UX

Linux

Solaris

Note: The **man** command prompt option is available in the AIX, HP-UX, Linux, and Solaris command prompt environments.

Windows Server

Windows Server users

Note: This feature is not available in the Windows Server command prompt environment.

The available man page (manual page) options are:

```
man xiv_attach
man xiv_fc_admin
man xiv_iscsi_admin
man xiv_syslist
man xiv_devlist
man xiv_diag
man xiv_host_profiler
```

Each of the man page options above displays content that is unique to the specified HAK utility.

Chapter 4. Single host attachment

Refer to this information when you are ready to attach your host to the IBM storage system.

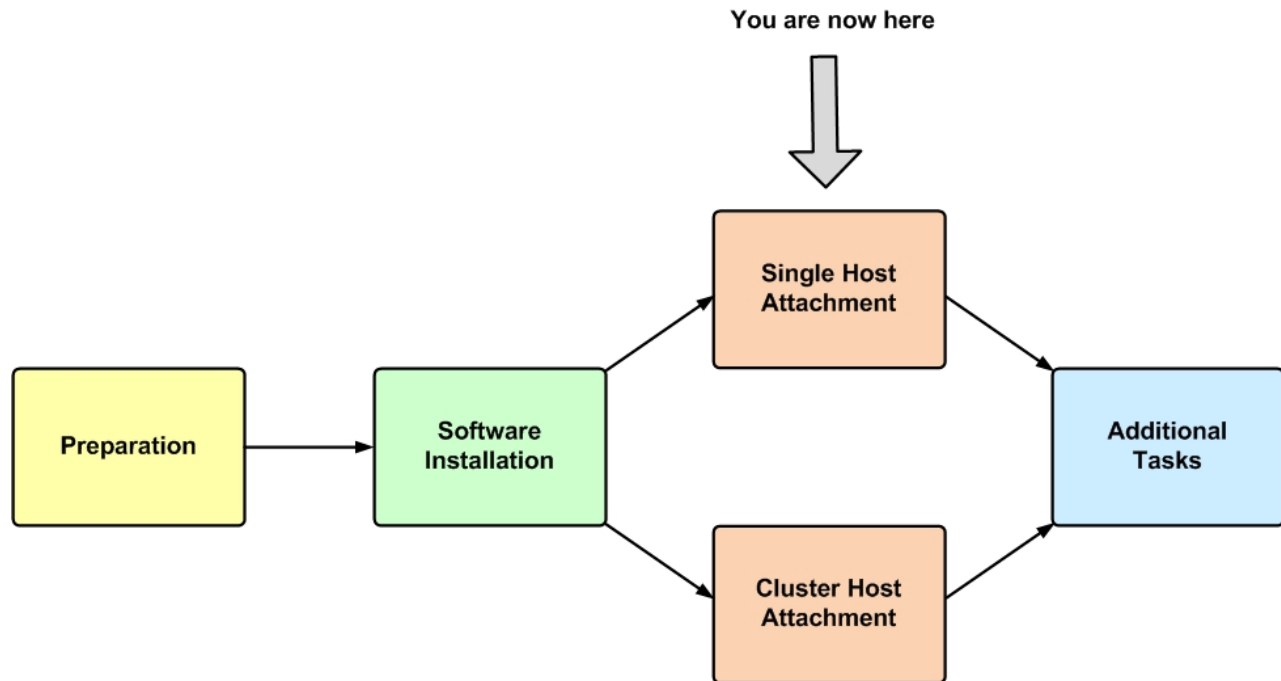


Figure 22. Single host attachment

The attachment of a single host includes the following stages:

- Stage 1: Using the CLI attachment utility (`xiv_attach`)
- Stage 2: Defining the host on the storage system (if applicable)
- Stage 3: Mapping storage volumes to the host

Stage 1: Using the CLI attachment utility (`xiv_attach`)

The IBM Storage Host Attachment Kit (HAK) provides an interactive command-line utility, `xiv_attach`, which locally configures and defines the host on the IBM storage system and rescans the already-attached storage device for new iSCSI ports.

About this task

Running the `xiv_attach` command initiates an interactive CLI wizard, which configures and defines the host on the IBM storage system. In addition, you can use the `xiv_attach` command to scan already-attached storage systems for new iSCSI ports (see “Issuing the `xiv_attach` command to connect to new targets” on page 43).

Procedure

When you are ready to start the host attachment, complete the following procedure.

1. Start the command-prompt program on the host.

Windows Server	Windows Server users Important: To bypass the account control limitations, start the program as follows: go to C:\WINDOWS\System32, right-click cmd.exe , and then select Run as Administrator .
Linux	Linux users Note: In the RHEL 7.x environment, you can enable the host-side acceleration (HSA) for improving iSCSI communication between Spectrum Accelerate and a host. The HSA uses <i>dm-switch</i> , a Linux module for device mapping. As a result, the HSA cannot be enabled, if a third-party multipathing framework is installed and activated on the host.

2. Enter **xiv_attach**. The Host Attachment Wizard welcome message is displayed.

```
Welcome to the IBM Host Attachment wizard, version 2.6.0.
This wizard will assist you to attach this host to the storage system.

The wizard will now validate host configuration for the storage system.
Press [ENTER] to proceed.
```

Note: You can run **xiv_attach** in verbosity mode using the **-i** or **--verbose** argument, as explained in “Displaying information when working with xiv_attach” on page 80.

3. Press **Enter**. The connectivity type choice message is displayed.

```
Please choose a connectivity type, [f]c / [i]scsi :
```

4. Enter the first letter of the connectivity type that the host uses to connect to the IBM storage system f for Fibre Channel or i for iSCSI.

- If you choose iSCSI and the host operating system supports the HSA, the following message is displayed:

```
Would you like to enable host-side acceleration? (A scheduled task will be created)
[default: no]: y
```

- If you enable the HSA, the following verification results are displayed:

```
Please wait while the wizard validates your existing configuration...
```

```

Verifying dm-switch kernel module installation      OK
Verifying dm-switch cron job                      OK
Verifying dm-switch systemd service              OK
Verifying multipath - multipath.conf              OK
Verifying multipath service(s)                   OK
Verifying iSCSI initiator name - initiatorname.iscsi  OK
Verifying iSCSI daemon - iscsid.conf             OK
Verifying iSCSI service                          OK

```

- If you do not use the HSA, the following verification results are displayed (in this Linux-specific example, all verification results are OK):

```
Please wait while the wizard validates your existing configuration...
```

```

Verifying multipath - multipath.conf              OK
Verifying multipath service(s)                   OK
Verifying iSCSI initiator name - initiatorname.iscsi  OK
Verifying iSCSI daemon - iscsid.conf             OK
Verifying iSCSI service                          OK

```

Note: If the host has an active third-party multipath framework (Linux only), the wizard displays the following message.

```
Warning: Non-native multipath framework has been detected on this host.
Would you like to use it instead of the native multipath framework? [default: no]: y
```

Select *y* if you want to use the third-party multipath framework, or *n* to utilize the default Linux multipathing device. This setting has an impact on the multipath mode (*-m*) attribute of the `xiv_iscsi_admin -m` or `xiv_fc_admin -m` command:

- Enabling a non-native multipath framework (*y* setting), sets the multipath mode to *IGNORE*.
 - Enabling a native multipath framework (*n* setting), sets the multipath mode to *DEFAULT*.
-

- If you choose FC, the following verification results are displayed (in this Linux-specific example, all verification results are OK):

```

Verifying multipath configuration - multipath.conf      OK
Verifying multipath service(s)                         OK

```

Important:

- Depending on the host configuration and installed components, you might be requested to restart the host. In such a case, restart the host and then run **xiv_attach** again.
- If the result of any verification is 'NOT OK', the Host Attachment Wizard requests your permission to continue:

The wizard needs to configure this host for the IBM storage system.
Do you want to proceed? [default: yes]:

The wizard then automatically performs the configuration and notifies you about any task that is being performed. For example, when the multipath services are being configured, the following prompt is displayed:

Configuring multipath service(s)...

Linux**Linux users**

Note: When you run **xiv_attach**, the Linux `multipath.conf` file is automatically modified. Although this modification was fully tested and is not known to cause conflicts or errors, the original file is automatically backed up to the `/backup` folder in the Host Attachment Kit directory (default: `/opt/xiv/host_attach/backup`). If the `multipath.conf` file does not exist on the host, the wizard creates it.

Solaris**Solaris users**

Important: The wizard scans for existing dynamic multipathing solutions. Valid solutions for the Solaris operating system are Solaris Multiplexed I/O (MPxIO) or Veritas Dynamic Multipathing (VxDMP).

If VxDMP is already installed and configured on the host, it is preferred over MPxIO and the following message is displayed:

Notice: VxDMP is available and
will be used as the DMP software.

If VxDMP is installed but cannot be used, the wizard prompts you to use MPxIO and the following message is displayed:

Warning: VxDMP is available
but is not usable: VxDMP is not enabled.
Would you like to proceed to use
MPxIO multipathing? [default: no]:

Windows Server

Windows Server users

Note: Windows Server hotfixes are automatically installed upon the first use of **xiv_attach** on the host. For more information about the automatically installed hotfixes (Microsoft KBs), refer to the IBM Storage Host Attachment Kit release notes for Windows.

5. After the host information is collected and analyzed, and if no scheduled task for the Host Profiler utility is defined, the following message is displayed:

```
Creating a host profile enables focused and better support for your host.  
Would you like to add a scheduled task for host profile creation? [default: yes ]:
```

Enter **y** to create a scheduled task for the Host Profiler utility. For more information about the Host Profiler utility, see “Using the Host Profiler utility” on page 85. The following message is displayed and the scheduling is set for each week on Sunday, at 4:00 AM:

```
A scheduled task for xiv_host_profiler has been added.
```

Note: If a scheduled task for the Host Profiler utility is already defined on the host, this step does not appear. If you remove the scheduled task (as explained in Table 14 on page 86) and then run **xiv_attach** again, this step reappears in the attachment procedure.

6. If you selected iSCSI and the IBM Storage Host Attachment Kit is running on Windows or Linux operating systems, the following message is displayed:

```
Discovering new iSCSI targets ...
```

Any new iSCSI ports for each already-connected storage system are discovered and connected to, automatically. For Windows operating systems, you might be requested to manually enter the CHAP credentials, depending on the host configuration.

7. Based on the chosen connectivity type (see step 3 on page 38), the wizard offers to attach the host to new storage arrays as follows:

- For Fibre Channel (FC) connectivity, the wizard displays the worldwide port names (WWPNs) of all Fibre Channel HBAs that you can zone. Press Enter to proceed.

```
Please zone this host and add its WWPNs with the IBM storage system:  
21:00:00:24:ff:25:4f:fc: [QLOGIC]: N/A  
21:00:00:24:ff:25:4f:fd: [QLOGIC]: N/A  
Press [ENTER] to proceed.
```

- For iSCSI connectivity, you are asked if you want to discover new iSCSI targets in order to connect to a new storage system. Enter **yes** (or **y**) to start. You are then asked to enter the iSCSI IP addresses and CHAP authentication details if these details are predefined for the host on the IBM storage system.

```

Would you like to discover a new iSCSI target? [default: yes]:
Please enter an iSCSI discovery address (iSCSI interface):
Is this host defined to use CHAP authentication with the IBM storage system?
[default: no]:
Enter CHAP name:
Enter CHAP secret:

```

- After the connectivity parameters are set, the wizard prompts you to start scanning for storage devices. Enter **yes** (or **y**) to start the scanning.

```

Would you like to rescan for new storage devices now? [default: yes]:

```

The following message is displayed:

```

Please wait while rescanning for storage devices...

```

The host attachment wizard scans for physically connected IBM storage systems and displays a list of the detected devices (arrays), as shown in the following example:

```

This host is connected to the following IBM storage arrays:
Storage Type  Serial  System Version  Host Defined  Ports Defined  Protocol  Host Name(s)
FlashSystem  MN67003  12.0.0          Yes           All            iSCSI    rhe17.1-64-hak
A9000R
XIV           MN65024  11.3.1.c       Yes           All            iSCSI    rhe17.1-64-hak
XIV           MN65022  10.2.4.e-8     Yes           All            iSCSI    rhe17.1-hak

```

- If the host is already defined on all the detected IBM storage systems, press Enter to exit the host attachment wizard.

```

This host is defined on all FC-attached IBM storage systems.
Press [ENTER] to proceed.
The IBM host attachment wizard has successfully configured this host.
Press [ENTER] to exit.

```

- If the host is not defined on all the detected IBM storage systems, you are asked whether you want to define the host on these systems. Enter the requested details as prompted:

```

Do you want to define this host on these systems now? [default: yes ]: yes
Please enter a name for this host [default: localhost.localdomain ]: <host name>
Please enter a username for system MN00022 : [default: admin ]: <admin user name>
Please enter the password of user admin for system MN00022:
Connecting to the storage array...
Press [ENTER] to proceed.
The IBM host attachment wizard has successfully configured this host.
Press [ENTER] to exit.

```

Note: If you are attaching the host to a Spectrum Accelerate, FlashSystem A9000/A9000R or an XIV Storage System of microcode 11.5.0 or later, you might be attaching the host to a specific domain or domains on that storage system, and not to the entire device. Attaching is determined by the host definition on the storage system, or by the credentials that you provide on the last attachment steps (see step 8).

Table below details the codes returned by the **xiv_attach** command to specify the execution status.

Table 3. `xiv_attach` exit codes

Code	Description
0	Successful execution
1	General error
2	Command executed with wrong arguments
3	Successful execution, but the host must be rebooted
4	Logging error
5	Verification stage failed
6	Command execution aborted by user

Issuing the `xiv_attach` command to connect to new targets

When an interface module is added to a storage system, a connected host needs to update its connections to the storage system. You can scan already-attached storage systems for new iSCSI ports by running the `xiv_attach` command. The `xiv_attach` command discovers all the new iSCSI targets and connects to them automatically. This process is best for cases where an interface (iSCSI) module is added or replaced.

Stage 2: Defining the host on the storage system (if applicable)

If needed (see the following notes), you can manually define the host on the IBM storage system.

AIX	<p>AIX users</p> <p>Important: This stage is mandatory for non-updated AIX versions that do not support automatic host definition on the IBM storage system. For more information about the required AIX updates, refer to the release notes. However, if your host uses the most updated AIX version, this stage is optional.</p>
HP-UX	<p>HP-UX, Linux, Solaris and Windows users</p> <p>Note: This stage is optional for HP-UX, Linux, Solaris, and Windows Server hosts.</p> <p>If the <code>xiv_attach</code> automatic host definitions were sufficient for your uses, skip to “Stage 3: Mapping storage volumes to the host” on page 49.</p>
Linux	
Solaris	
Windows Server	

Depending on your satisfaction (and the host operating system; AIX or other) from the automatic host configuration performed by the `xiv_attach` utility (see “Stage 1: Using the CLI attachment utility (`xiv_attach`)” on page 37), you can manually define the host on the storage system using the management GUI or the CLI, as described in the following subsections:

- “Defining a host from the XIV GUI” on page 44
- “Defining a host from the CLI” on page 45

- “Adding ports to a host from the XIV GUI” on page 46
- “Defining a host from the IBM Hyper-Scale Manager” on page 46

Defining a host from the XIV GUI

This section applies if you want to use the GUI for defining a host.

Procedure

Perform the following procedure to configure a host from the GUI.

1. Log in to the GUI management application and zoom into the XIV Storage System that you want to configure.
2. Move the mouse over the **Hosts and Clusters** icon and then click **Hosts and Clusters**.



Figure 23. XIV GUI - Hosts and Clusters

The Hosts and Clusters table is displayed.



Figure 24. XIV GUI - Hosts and Clusters table

Note: If the host already appears on the list, you can right-click its name and then select **Edit** to edit its details.

3. Click the **Add Host** button located at the top of the window. The Add Host dialog box is displayed.



Figure 25. XIV GUI - Add host dialog box

4. Enter the following host details:
 - **Name** – Name of the host.
 - **Cluster** – If the host is a member of a cluster, select the cluster name from the drop down list. Otherwise, leave the **None** default selection.
 - **Type** – Select the host type from the drop down list.
 - **CHAP Name** – If you want the host to use CHAP authentication over iSCSI for this host (not applicable for AIX), enter the CHAP name. Otherwise, leave this field blank.
 - **CHAP Secret** – If a CHAP name was entered above, enter the CHAP secret. Otherwise, leave this field blank.
5. Click **Add**. The host definition is added to the XIV storage system.

Defining a host from the CLI

If you want to define a host from the CLI rather than from the GUI, use the following commands.

- To list a specific host or all hosts:

```
host_list [host=HostName]
```

- To define a new host and connect it to the IBM storage system:

```
host_define host=HostName [cluster=ClusterName]
```

For more detailed information about host definition CLI commands, refer to:

- IBM XIV XCLI Reference Guide, available in the IBM XIV Storage System documentation (www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).
- IBM Spectrum Accelerate™ Command-Line Interface (CLI) Reference Guide, available in the IBM Spectrum Accelerate documentation (http://www.ibm.com/support/knowledgecenter/STZSWD/landing/IBM_Spectrum_Accelerate_welcome_page.html).
- IBM FlashSystem A9000 CLI Reference Guide, available in the IBM FlashSystem A9000 documentation (www.ibm.com/support/knowledgecenter/STJKMM).

- IBM FlashSystem A9000R CLI Reference Guide, available in the IBM FlashSystem A9000R documentation (www.ibm.com/support/knowledgecenter/STJKN5).

Adding ports to a host from the XIV GUI

After the host is defined on the IBM storage system, it requires port definitions.

Procedure

If no ports have been defined for the host, perform the following procedure to add iSCSI or FC ports.

1. On the Hosts and Clusters table (see Figure 24 on page 44), right-click the host and select **Add Port**. The Add Port dialog box is displayed.



Figure 26. XIV GUI - Add Port dialog box

2. From the **Port Type** drop down list, select **FC** or **iSCSI** to define the port type.
3. Depending on the selected port type, set the port name as follows:
 - For an FC port – Select the World Wide Port Name (WWPN) of the port from the drop down list.
 - For an iSCSI port – Enter the iSCSI Qualified Name (IQN) of the port.
4. Click **Add**. The port definition is added to the host.

Defining a host from the IBM Hyper-Scale Manager

This section applies if you want to use the IBM Hyper-Scale Manager management application for defining a host.

Procedure

Perform the following procedure to configure a host from the Hyper-Scale Manager application.

1. Log in to the IBM Hyper-Scale Manager application.
2. Click the **New** icon on the toolbar at the top and then select **Host**.

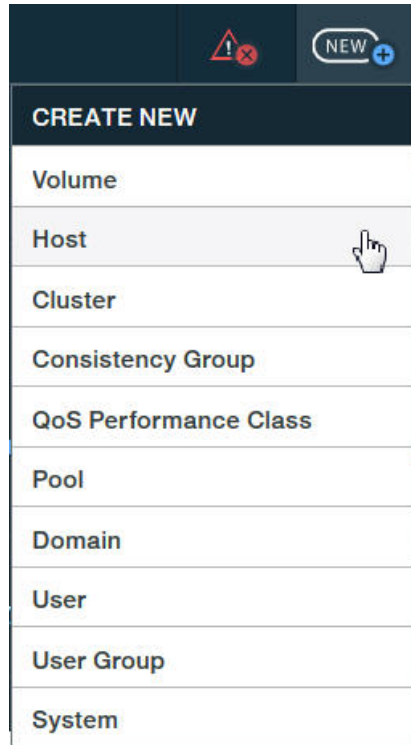


Figure 27. New Host option

The **Add Host** form is displayed.

The screenshot shows the 'Add Host' configuration form. At the top right, there is a blue 'NEW' badge. The form title 'Add Host' is on the left. Below the title, there are two main input fields: 'Name' with a red asterisk indicating it is required, and 'Type' with a dropdown menu currently set to 'Default'. A section titled 'PORTS' has a plus sign to its right. Inside this section, there are two radio buttons: 'FC' (selected) and 'iSCSI'. To the right of the radio buttons is a 'Port Address' dropdown menu with a red asterisk. Below the 'PORTS' section is another section titled 'SYSTEMS / CLUSTERS / DOMAINS' with a plus sign to its right. This section contains two dropdown menus: 'System' and 'Cluster', both with red asterisks. Below these are two more sections: 'DOMAINS' and 'CLUSTER'S DOMAINS', each with a plus sign. At the bottom right of the form are two buttons: 'Cancel' and 'Apply'.

Figure 28. Add Host form

3. Enter the following host details:
 - **Name** – Define name of the host.
 - **Type** – Select the host type from the drop down list.
 - **PORTS** – Select **FC** or **iSCSI** to define the port type. Depending on the selected port type, set the port address as follows:
 - For an FC port – Select the World Wide Port Name (WWPN) of the port from the drop down list.
 - For an iSCSI port – Enter the iSCSI Qualified Name (IQN) of the port.
 - **System** – Select the storage system to which you want to attach the host.
 - **Cluster** – If the host is a member of a cluster, select the cluster name from the drop down list. Otherwise, leave this option blank.
 - **Domain** – Select the storage system domain or domains in which this host should be defined.
 - **Cluster's Domain** – Select the cluster domain or domains in which this host should be defined.
4. Click **Apply**. The host definition is added to the storage system.

Stage 3: Mapping storage volumes to the host

After the `xiv_attach` utility has detected the storage systems, and after the host has been defined on these storage systems, you can use the management GUI to map volumes (LUNs) to the host.

Use either the XIV GUI or the IBM Hyper-Scale Manager application to map one or more volumes to the host, as explained in the following sections:

- “Mapping storage volumes to the host from the XIV GUI”
- “Mapping storage volumes to the host from the IBM Hyper-Scale Manager” on page 50

Mapping storage volumes to the host from the XIV GUI

This section applies if you want to use the XIV GUI for mapping a volume to a host.

Procedure

Perform the following procedure to map one or more volumes to the host, using the XIV GUI.

1. On the Hosts and Clusters table (see Figure 24 on page 44), right-click the host to which you want to map volumes, and then select **Modify LUN Mapping**.

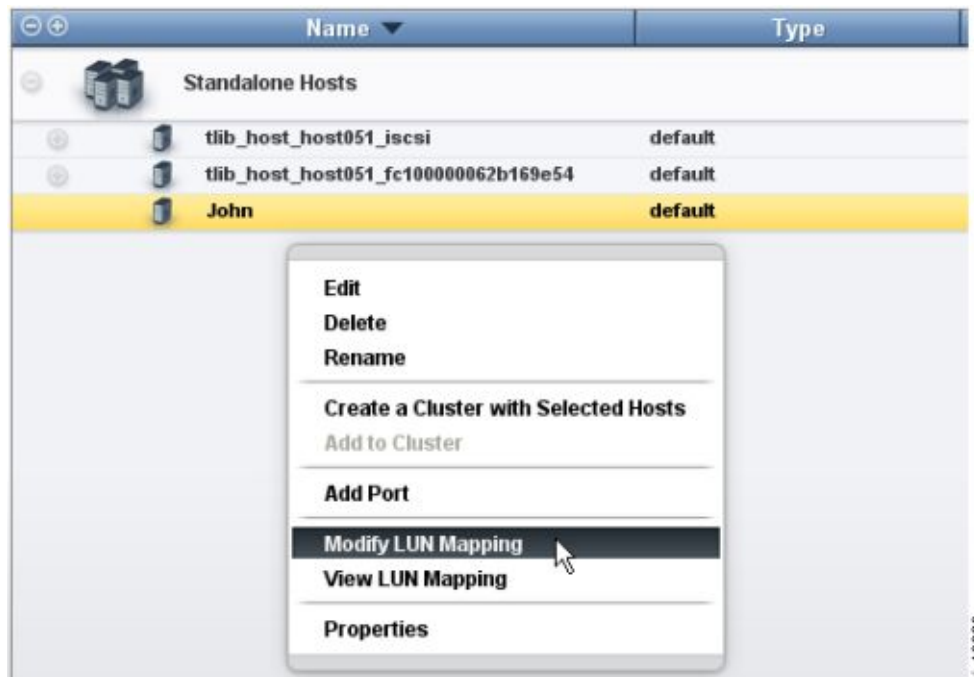


Figure 29. XIV GUI - Modifying LUN Mapping to a host

The LUN mapping table is displayed.

2. Select one or more volumes that you want to map to the host, and then click **Map**.




Figure 30. XIV GUI - Modifying LUN Mapping

The volume mapping is complete.

Mapping storage volumes to the host from the IBM Hyper-Scale Manager

This section applies if you want to use the IBM Hyper-Scale Manager management application for mapping a volume to a host.

Procedure

Perform the following procedure to map one or more volumes to the host, using the IBM Hyper-Scale Manager application.

1. Log in to the IBM Hyper-Scale Manager application.
2. On the sidebar, click the **Pools and Volumes Views** icon, and then click **Volumes** to display the list of existing volumes.
3. On the **Volume** table, select one or more volumes to which you want to map a host. The **Volume** mini-dashboard is displayed.



Figure 31. Volume mini-dashboard

4. From the **Volume** mini-dashboard, click the **Mapping** relation icon to display the empty **Volume Mapping** form with the **ADD** icon.
5. In the **Volume Mapping** form, click **ADD** icon to map the volume to an existing host.

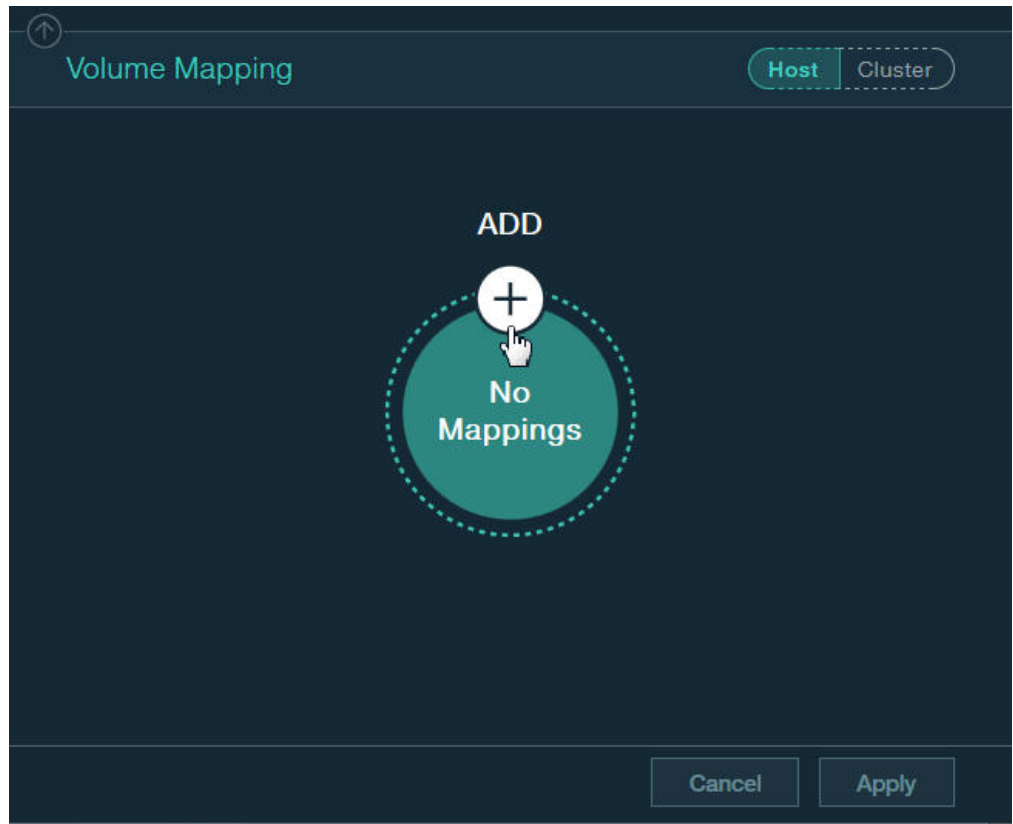


Figure 32. Adding new volume mapping to host

The **Host** and **LUN** drop-down lists are displayed in **Volume Mapping** form.




Figure 33. Adding new volume mapping for a host

6. In the **Volume Mapping** form, select a host and a LUN to be mapped, using the respective drop-down lists.
7. Click the plus sign in the right-hand corner if you want to continue the host mapping.
8. Click **Apply** to finish the mapping.

Chapter 5. Cluster host attachment

Refer to this information if you want to attach a cluster of hosts (as opposed to a single host) to the IBM storage system.

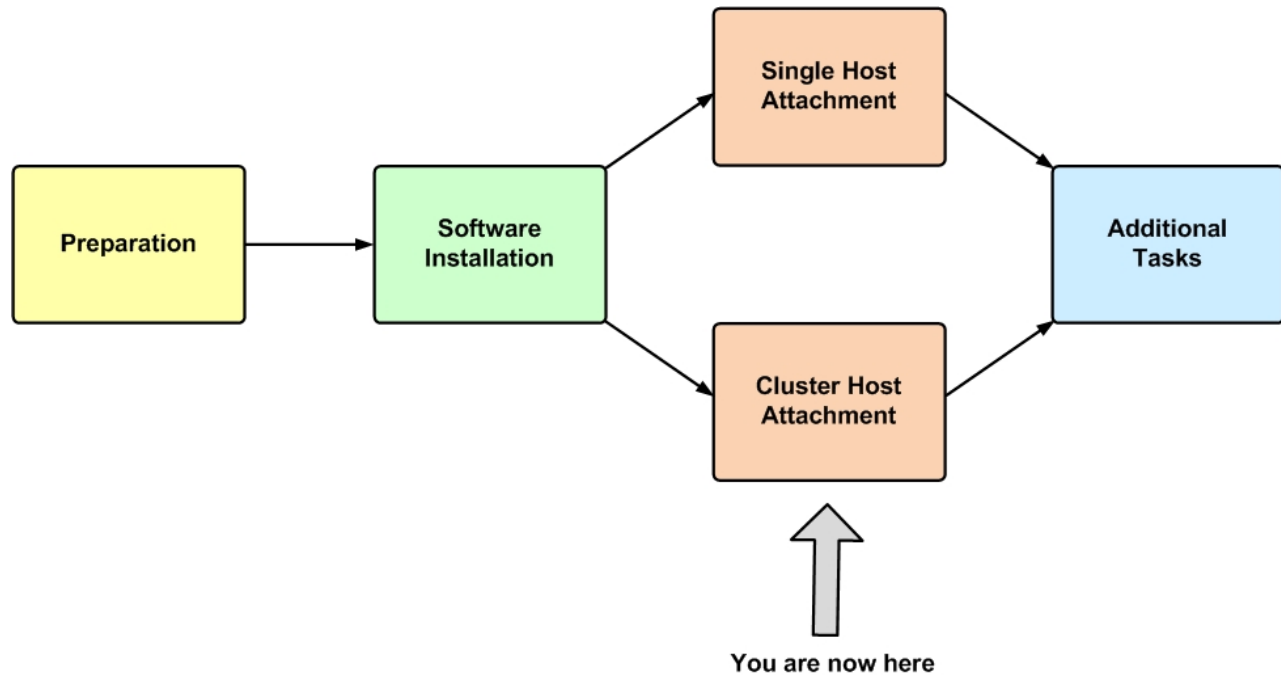


Figure 34. Cluster host attachment

Attaching a group of hosts (see Figure 35 on page 54) requires different prerequisites and a different host attachment procedure. This includes:

- Cluster installation prerequisites
- Stage 1: Defining the cluster on the storage system
- Stage 2: Mapping storage volumes to the cluster
- Stage 3: Running `xiv_attach` on each cluster node




Figure 35. A cluster connected to the IBM storage systems

Cluster installation prerequisites

Install or mount the IBM Storage Host Attachment Kit (HAK) software on all cluster nodes. If you are creating a new cluster, you can install the HAK simultaneously.

Procedure

If you are attaching an existing cluster to an IBM storage system, install the HAK on the cluster nodes in the following manner:

1. Install the HAK on the passive nodes first.
2. Perform failover of the cluster groups and resources from the active nodes to the passive nodes on which you have installed the HAK.
3. Install the HAK on the remaining nodes.

Stage 1: Defining the cluster on the storage system

All hosts that belong to a cluster must be defined on the IBM storage system side as cluster members.

Use either the GUI or the CLI method to define cluster hosts, as explained in the following sections:

- “Defining a cluster from the XIV GUI” on page 55

- “Defining a cluster from the CLI” on page 57
- “Defining a cluster from the IBM Hyper-Scale Manager” on page 57

Defining a cluster from the XIV GUI

This section describes how to define a cluster from the GUI.

Procedure

1. Log in to the IBM XIV GUI management application and zoom into the XIV storage system that you want to configure.
2. Move the mouse over the Hosts and Clusters icon and then click **Hosts and Clusters** (see Figure 23 on page 44). The Hosts and Clusters table appears (see Figure 24 on page 44).
3. Click the **Add Cluster** button located at the top of the window. The Add Cluster dialog box is displayed.



Figure 36. Add Cluster dialog box

4. Enter the following cluster details:
 - **Name** – Name of the cluster.
 - **Type** – Select the cluster type from the drop-down list.
5. Click **Add**. The cluster is added to the Hosts and Clusters table.



Figure 37. XIV GUI - Cluster added to the Hosts and Clusters table

6. On the Standalone Hosts list, right-click a host that you want to add to the cluster, and then select **Add to Cluster**.




Figure 38. XIV GUI - Add to Cluster

Note: You can select more than one host by pressing **CTRL** on your keyboard and clicking the hosts that you want to add. After selecting multiple hosts, right-click the **yellow selection area** to add all selected hosts to the cluster.

The Add Host(s) to Cluster dialog box appears.



Figure 39. XIV GUI - Add Host(s) to Cluster dialog box

- From the drop-down list, select the name of the cluster to which you want to add the selected host(s), and then click **Add**. All selected hosts are added to the cluster.

Name	Type	Cluster	Access
Standalone Hosts			
tlib_host_host051_iscsi	default		
John			
tlib_host_host051_fc10000062b169e54	default	John	default

Figure 40. XIV GUI - Host added to a cluster

Defining a cluster from the CLI

If you want to define a cluster from the CLI rather than from the GUI, use the following commands.

- To create a new cluster:

```
cluster_create cluster=ClusterName
```

- To add an existing host to a cluster:

```
cluster_add_host cluster=ClusterName host=HostName map=<cluster|host>
```

- To list existing clusters:

```
cluster_list [cluster=ClusterName]
```

For more detailed information about cluster definition CLI commands, refer to:

- IBM XIV XCLI Reference Guide, available in the IBM XIV Storage System documentation(www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).
- IBM Spectrum Accelerate Command-Line Interface (CLI) Reference Guide, available in the IBM Spectrum Accelerate documentation(http://www.ibm.com/support/knowledgecenter/STZSWD/landing/IBM_Spectrum_Accelerate_welcome_page.html).

Defining a cluster from the IBM Hyper-Scale Manager

This section applies if you want to use the IBM Hyper-Scale Manager management application for defining a cluster.

Procedure

Perform the following procedure to configure a cluster from the Hyper-Scale Manager application.

- Log in to the IBM Hyper-Scale Manager application.
- Click the **New** icon on the toolbar at the top and then select **Cluster**.

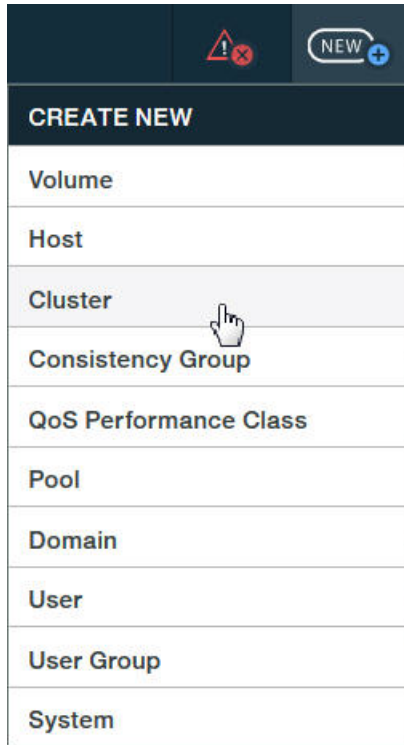


Figure 41. New Cluster option

The **Create Cluster** form is displayed.

Figure 42. Create Cluster form

3. Enter the following host details:
 - **Name** – Define name of the cluster.
 - **Type** – Select the cluster type from the drop down list.
 - **System** – Select the storage system on which you want to define the cluster.
 - **Domain** – Select the storage system domain or domains in which this cluster should be defined.
4. Click **Apply**. The cluster definition is added to the storage system.
5. Add a single or multiple hosts to the cluster, as explained in (see “Defining a host from the IBM Hyper-Scale Manager” on page 46) The host and the cluster become listed in the **Hosts & Clusters** table.

HOST ^	TYPE	FC PO...	ISCSI ...	CHAP ...	SYSTEM	DOMAINS	CLUSTER
asd	Default	0	0		2	4	gfn

Figure 43. Hosts & Clusters table

Stage 2: Mapping storage volumes to the cluster

Storage volumes must be manually mapped to the cluster.

Two different volume types can be mapped to a cluster:

- **Shared volumes** – Mapped to all the nodes (hosts) in the cluster. Shared volumes are typically used in clustered application environments.

- **Private volumes** – Mapped to a specific node (host) in the cluster. Private volumes are required in environments in which a specific cluster node requires exclusive access to a volume. For example, in environments that utilize boot volumes, each cluster node requires exclusive access to a specific boot volume, so that other cluster nodes would not corrupt the operating system on that boot volume. In these environments, map a different private boot volume to each host, and map the shared volumes to the entire cluster.

Depending on the volume type that you want to map to the cluster, refer to one of the following sections:

- “Mapping shared volumes to all the hosts in a cluster from the XIV GUI”
- “Mapping private volumes to a specific host in a cluster from the XIV GUI” on page 61
- “Mapping storage volumes to the cluster from the IBM Hyper-Scale Manager” on page 62

Mapping shared volumes to all the hosts in a cluster from the XIV GUI

This section describes how to map shared volumes to all the hosts in a cluster from the XIV GUI.

Procedure

1. On the Hosts and Clusters table (see Figure 37 on page 55), right-click the cluster to which you want to map volumes, and then select **Modify LUN Mapping**.



Figure 44. XIV GUI - Modify LUN Mapping to a cluster

2. Select one or more volumes that you want to map to the cluster, and then click **Map**.

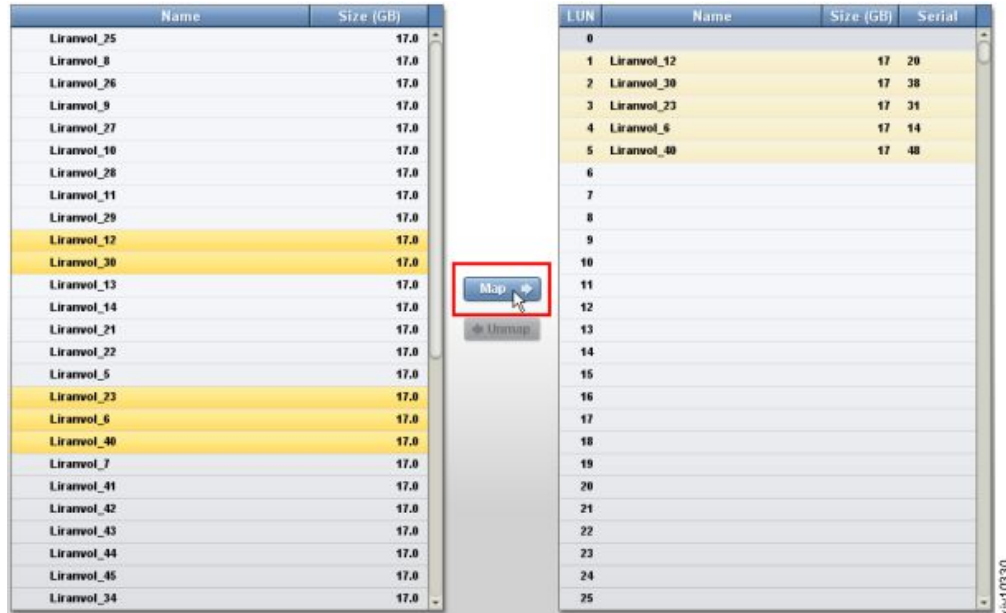


Figure 45. XIV GUI - Modify LUN Mapping

The volume mapping is complete.

Mapping private volumes to a specific host in a cluster from the XIV GUI

This section describes how to map private volumes to a specific host in a cluster from the XIV GUI.

To map a volume exclusively to a specific host of a cluster, right-click the host name on the list of clustered hosts.




Figure 46. Right-clicking a specific host of a cluster

The rest of the private volume mapping is performed in the same manner as in shared volume mapping (see Figure 45 on page 61).

Note: After mapping private volumes to a host of a cluster, the private mapping is labeled in the mapped volumes list as **Private Mapping**.

Mapping storage volumes to the cluster from the IBM Hyper-Scale Manager

This section applies if you want to use the IBM Hyper-Scale Manager management application for mapping a volume to a cluster.

Procedure

Perform the following procedure to map one or more volumes to the cluster, using the IBM Hyper-Scale Manager application.

1. Log in to the IBM Hyper-Scale Manager application.
2. On the sidebar, click the **Pools and Volumes Views** icon, and then click **Volumes** to display the list of existing volumes.
3. On the **Volume** table, select one or more volumes to which you want to map a cluster. The **Volume** mini-dashboard is displayed.



Figure 47. Volume mini-dashboard

4. From the **Volume** mini-dashboard, click the **Mapping** relation icon to display the empty **Volume Mapping** form with the **ADD** icon.
5. In the **Volume Mapping** form, click the **Cluster** button, and then click **ADD** icon to map the volume to an existing cluster.

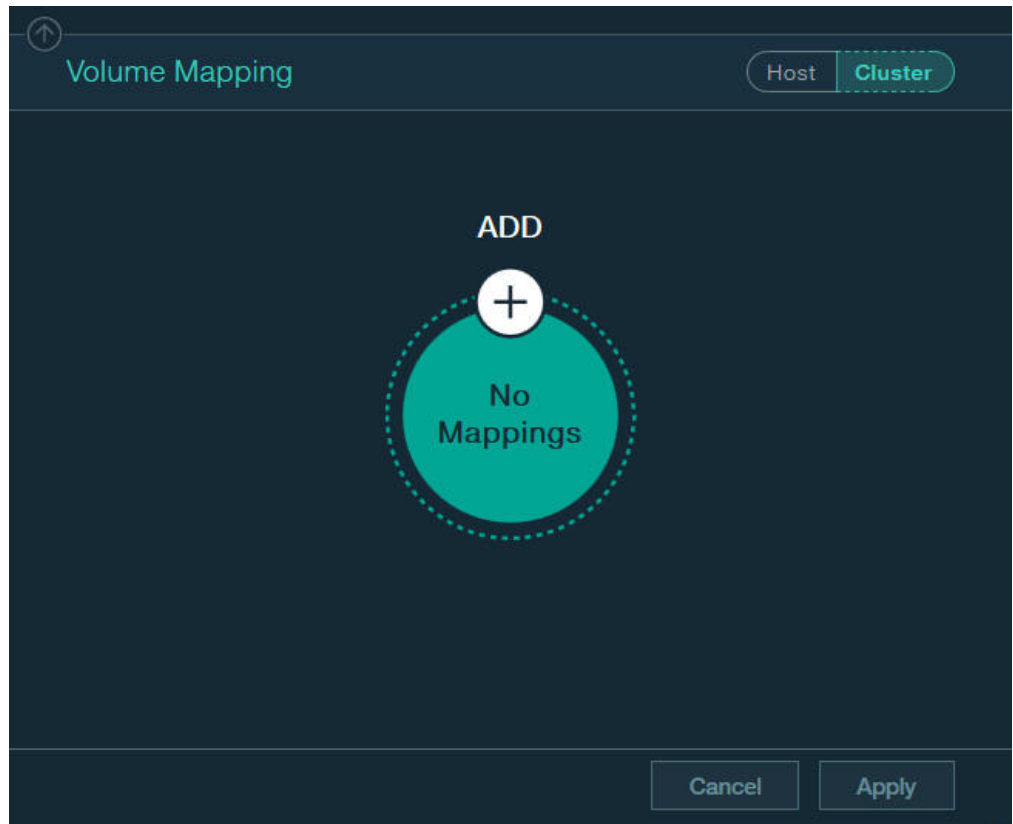


Figure 48. Adding new volume mapping to cluster

6. In the **Volume Mapping** form, click .
7. In the **Volume Mapping** form, select a cluster and a LUN to be mapped, using the respective drop-down lists. The **Host** and **LUN** drop-down lists are displayed in **Volume Mapping** form.

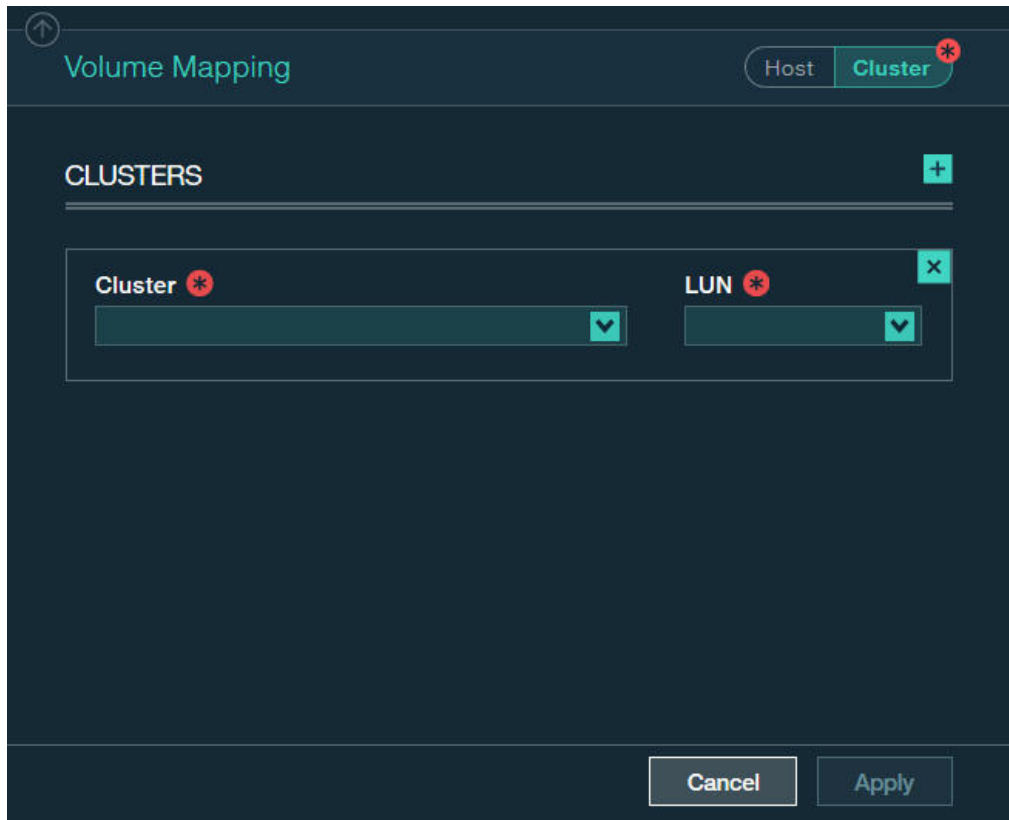


Figure 49. Adding new volume mapping for a cluster

8. Click the plus sign in the right-hand corner if you want to continue the cluster mapping.
9. Click **Apply** to finish the mapping.

Stage 3: Running `xiv_attach` on each cluster node

After the cluster is defined and storage volumes are mapped to the cluster on the IBM storage system side, run the `xiv_attach` utility on each and every node in the cluster.

For more information about how to use the `xiv_attach` utility, see Chapter 4, "Single host attachment," on page 37 (previous chapter; refer only to Stage 1).

Chapter 6. Additional tasks

Refer to this information to learn about useful tools that are included in the IBM Storage Host Attachment Kit (HAK).




Figure 50. Additional tasks

After your host or cluster has been successfully attached to the IBM storage system, you can perform different administrative operations on the host by using the different HAK utilities.

These operations include the following optional tasks:

- Displaying information about connected storage systems
- Displaying volume and additional information
- Setting Fibre Channel connectivity parameters
- Setting iSCSI connectivity parameters
- Displaying information about host configuration changes – applicable only to Linux and Windows Server

To learn about diagnostic tasks, see Chapter 7, “Diagnostics,” on page 83.

AIX

AIX users

Important: To avoid host reboots and downtime when using the FC configuration command (`xiv_fc_admin -C`), it is recommended to always use the latest AIX TL update and service pack (SP) on the host. The latest AIX updates have native support for the storage system. For more information, refer to the release notes.

Displaying information about connected storage systems

If you want to display the details of all IBM storage systems that are currently detected as physically connected to the host, use the **xiv_syslist** command.

The listed storage systems may be ones to which the host is already attached, or ones to which the host is not yet attached (for example, if new systems were added recently).

Note: Storage systems to which the host is not yet attached are detected by **xiv_syslist** as long as the FC or iSCSI connectivity was properly set. For more information about how to properly set the FC or iSCSI connectivity between the host and the storage system, refer to Chapter 2, "Preparation," on page 5.

The following example shows the default **xiv_syslist** output format for two XIV systems that are connected to the host:

```
xiv_syslist

This host is connected to the following XIV storage arrays:
System Name: XIV hostdev32a
-----
System Details:
-----
System Serial Number: MN65024
System Version: 11.3.1.c
Management IPv4 Addresses: 9.151.156.153, 9.151.157.185
Management IPv6 Addresses:
Storage Type: XIV

-----
Connectivity:
-----
Connectivity Type: iSCSI
System iSCSI Targets: Module 4, Module 7

-----
Host Definition:
-----
Host Defined: Yes
Host Name Defined: alina_rh72
Number of Host Ports Defined: All

System Name: XIV Gen4d-67c
-----
System Details:
-----
System Serial Number: MN67003
System Version: 12.0.0-internal_67-p20160411_163317
Management IPv4 Addresses: 9.151.156.112
Management IPv6 Addresses:
Storage Type: FlashSystem A9000R

-----
Connectivity:
-----
Connectivity Type: iSCSI
System iSCSI Targets: Module 10, Module 9

-----
Host Definition:
-----
Host Defined: Yes
Host Name Defined: alina_rh72
Number of Host Ports Defined: All
```


The displayed details of any detected storage system include:

- System name – Name of the storage system.
- System details – System-specific details, including:
 - **System serial number** – Serial number of the storage system.
 - **System version** – Microcode version of the storage system.
 - **Management IPv4 addresses** (displayed if IPv4 is used) – The IPv4 address(es) through which the storage system is accessed and managed.
 - **Management IPv6 addresses** (displayed if IPv6 is used) – The IPv6 address(es) through which the storage system is accessed and managed.
 - **Storage Type** – Type of the storage system (XIV, Spectrum Accelerate, FlashSystem A9000 or FlashSystem A9000R).
- Connectivity – Connectivity information, including:
 - **Connectivity type** – Type of connectivity that is used to connect to the storage system: FC or iSCSI.
 - **System FC targets** (displayed if FC is used) – Lists the specific storage modules and FC port numbers with which FC connectivity is established.
 - **System iSCSI targets** (displayed if iSCSI is used) – Lists the specific modules with which iSCSI connectivity is established.
- Host definition – Information regarding this host as defined on the storage system, including:
 - **Host defined** – Indicates with 'Yes' or 'No' whether the host is defined on the storage system. If your host was properly attached to the storage system, the indication should be 'Yes'.
 - **Host name defined** (displayed if the host is defined) – Name of the host as defined on the storage system.
 - **Number of host ports defined** – Number of host ports that are defined on the storage system. If all host ports are defined, the output is 'All'.

If you want to display only specific storage system details in CSV or XML format, or if you want to save the information to a CSV or an XML file, add the required arguments after the **xiv_syslist** command, as detailed in Table 4.

Note: The customized output of **xiv_syslist** (using the **-o** argument, as detailed in the following table) is not available in regular screen view. You must specify either CSV or XML as the output format.

Important: All command arguments are case-sensitive.

Table 4. Optional arguments for **xiv_syslist**

Short-form argument	Long-form argument	Use after xiv_syslist to:
-t csv	--out csv	Display the xiv_syslist information on the screen (without saving to a file), in either CSV or XML format (as specified).
-t xml	--out xml	
		For the CVS output option, you can add the --no-headers option in the command line to generate the CSV output without column headers.

Table 4. Optional arguments for `xiv_syslist` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_syslist</code> to:
-L	--list	<p>Display the details of all IBM storage systems that are currently attached via Fibre Channel or iSCSI. These details include:</p> <ul style="list-style-type: none"> • Serial – Serial number of each storage system. • Version – Microcode version of each storage system. • Host defined – Indicates whether this host is defined on each of the listed storage systems. • Ports defined – Indicates which FC or iSCSI ports are defined on each of the listed storage systems. • Protocol – The communication protocol used with each storage system. • Host name – Name of this host as defined on each of the listed storage systems. • Storage type – Type of the storage system.
-f <file name> -t csv -t xml	--file <file name> --out csv --out xml	<p>Save the <code>xiv_syslist</code> information to a file.</p> <p>The name of the file must be typed after the -f or --file argument, and should be followed by the -t or --out argument that determines the file format: CSV or XML.</p> <p>For example:</p> <div style="border: 1px solid black; border-radius: 15px; padding: 5px; width: fit-content; margin: 10px auto;"> <pre>xiv_syslist -f <file name> -t csv</pre> </div>

Table 4. Optional arguments for `xiv_syslist` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_syslist</code> to:
<p>-t <csv or xml></p> <p>And then:</p> <p>-o system_name</p> <p>-o system_serial_number</p> <p>-o system_version</p> <p>-o management_ipv4_addresses</p> <p>-o management_ipv6_addresses</p> <p>-o storage_type</p> <p>-o host_defined</p> <p>-o number_of_host_ports_defined</p> <p>-o connectivity_type</p> <p>-o system_fc_targets</p> <p>-o system_iscsi_targets</p> <p>-o host_name_defined</p> <p>-o all</p>	<p>--out <csv or xml></p> <p>And then:</p> <p>--options</p> <p>(followed by a specific field)</p>	<p>Display specific details depending on the parameter(s) typed after the -o argument.</p> <p>Important: The -o or --options argument must follow the -t or --out argument (specifying CSV or XML).</p> <p>The specific details are as previously described in this topic, and are referred to in the command line as "fields":</p> <ul style="list-style-type: none"> • System name (system_name) • System serial number (system_serial_number) • System version (system_version) • Management IPv4 addresses (management_ipv4_addresses). • Management IPv6 addresses (management_ipv6_addresses). • Storage type (storage_type). • Host defined (host_defined) • Number of host ports defined (number_of_host_ports_defined) • Connectivity type (connectivity_type) • System FC targets (xiv_id) • System iSCSI targets (xiv_host) • Host name defined (host_name_defined) • All (all) – Use this parameter if you want to display all the details, including the ones that are not relevant and normally not displayed in the default on-screen <code>xiv_syslist</code> output. <p>If you want the customized output to include more than one specific detail, specify the fields after the -o argument, and use a comma without a space between each field name.</p> <p>For example:</p> <pre>xiv_syslist -t csv -o system_name,system_serial_number</pre>
-h	--help	Display usage format help information that is specific to <code>xiv_syslist</code> .

Table below details the codes returned by the `xiv_syslist` command to specify the execution status.

Table 5. `xiv_syslist` exit codes

Code	Description
0	Successful execution
1	General error
2	Command executed with wrong arguments
4	Logging error

Displaying volume and additional information

Use the `xiv_devlist` utility to display volume and additional information.

The following usage options are available:

- Displaying volumes and other storage devices
- Displaying additional information

Displaying volumes and other storage devices

You can use `xiv_devlist` to display the details of all volumes that are mapped to the host, as well as a list of other detected storage devices (volume mapping information for other devices is not available).

The following is an example of the displayed output when using `xiv_devlist`:

```
xiv_devlist
IBM storage devices
-----
Device          Size (GB) Paths  Vol Name          Vol ID Storage ID Storage Type      Hyper-Scale Mobility
-----
/dev/mapper/mpatha  34.4    24/24  il-bc01-b14-rhel62-hak  2668  MN00022  XIV              N/A
/dev/mapper/mpathai  17.2    6/6    yyyy_W_vol_002      321   9027529  Spectrum Accelerate Idle
/dev/mapper/mpathaj  17.2    6/6    yyyy_W_vol_001      320   9027529  Spectrum Accelerate Idle
-----
Non-XIV Devices
-----
Device  Size (GB) Paths
-----
/dev/sda  8.6      N/A
-----
```

The default `xiv_devlist` output displays the following information in columns:

- **Device** – Name of each detected volume, as defined at the operating system level.
- **Size (GB)** – Total capacity of each detected volume, in gigabytes. You can change the size unit with the `-u` argument, as explained in Table 6 on page 71
- **Paths** – Active vs. total available communication paths. If the host-side acceleration is enabled, the **Paths** column displays the total number of paths available for the Linux `dm-switch` device.
- **Volume name** – Name of each detected volume.
- **Volume ID** – ID number of each detected volume.
- **Storage ID** – ID of each detected storage device.
- **Storage Type** – Type of each detected storage device.
- **Hyper-Scale Mobility** – Current status of the Hyper-Scale Mobility volume migration.

Note: To refresh the displayed information, use the `xiv_fc_admin -R` or `xiv_iscsi_admin -R` command to rescan for storage devices. Then, run `xiv_devlist` again.

Displaying additional information

You can use **xiv_devlist** to display additional information by adding different optional arguments after the command.

The following table summarizes the different arguments and their use.

Important: All command arguments are case-sensitive.

Table 6. Optional arguments for **xiv_devlist**

Short-form argument	Long-form argument	Use after xiv_devlist to:
-V	--version	Display the HAK version number.
-x	--xiv-only	Display only the XIV-type systems, such as XIV, Spectrum Accelerate, FlashSystem 9000 or FlashSystem 9000R. Omit this argument to display all connected storage devices, XIV-type and non-XIV-type alike. See “Displaying volumes and other storage devices” on page 70.
-m auto -m native -m veritas	--multipath auto --multipath native --multipath veritas	Choose the multipath method by which the xiv_devlist information should be retrieved, depending on the parameter name typed after the -m or --multipath argument. The possible multipath methods are: <ul style="list-style-type: none">• Auto (auto) – automatically chosen.• Native (native) – native to the host operating system.• Veritas (veritas) – compatible with Symantec Veritas.

Table 6. Optional arguments for xiv_devlist (continued)

Short-form argument	Long-form argument	Use after xiv_devlist to:
-o device -o size -o serial -o lun -o multipath -o vendor -o vol_name -o vol_id -o xiv_id -o hyperscale_mobility -o xiv_host -o all	--options device --options size --options serial --options lun --options multipath --options vendor --options vol_name --options vol_id --options xiv_id --options hyperscale_mobility --options xiv_host --options all	<p>Display different information types depending on the parameter name typed after the -o or --options argument. The information types are:</p> <ul style="list-style-type: none"> • Device (device) – Name of each detected volume. • Size (size) – Total capacity of each detected volume. • Serial (serial) – Serial number of each detected volume. • LUN (lun) – Logical unit number of each detected volume. • Multipath (multipath) – Active vs. total available paths. • Vendor (vendor) – Vendor name of each detected storage device. • Volume name (vol_name) – Name of each detected volume. • Volume ID (vol_id) – ID number of each detected volume. • Storage ID (xiv_id) – ID of each detected storage system. • Hyper-Scale Mobility (hyperscale_mobility) – Current state of any IBM Hyper-Scale Mobility operation (online volume migration operation) for each detected volume. Applicable only to XIV microcode version 11.3.0 or later. The possible states are: <ul style="list-style-type: none"> – Idle – No volume migration operation is currently running. – In Process – A volume migration operation is currently in process and pending completion. – Proxy – The volume copy operation has started and the volume is currently mapped to both the source storage system and the target storage system. – Proxy Disabled – The migration operation is completed: the volume has been removed from the source storage system and is now mapped only to the target storage system. <p>Note: For more information about the IBM Hyper-Scale Mobility feature, refer to <i>IBM XIV Storage System – IBM Hyper-Scale Mobility Overview and Usage</i> (www.redbooks.ibm.com/redbooks.nsf/redpieceabstracts/redp5007.html).</p> <ul style="list-style-type: none"> • Storage host (xiv_host) – Name of this host as defined on each detected storage system. • All (all) – All information types. <p>Note: For devices other than IBM volumes, the Vendor column is also displayed when using the -o argument. The name of the relevant device vendor is displayed.</p> <p>You can type more than one parameter after -o (or --options) to display more than one information type. Use a comma to separate between the requested information types.</p> <p>For example:</p> <pre>xiv_devlist -o device,size,serial,lun</pre>
-l	--list	<p>Display all the information field types that are available for xiv_devlist --options.</p>

Table 6. Optional arguments for `xiv_devlist` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_devlist</code> to:
-u MB -u GB -u TB -u MiB -u GiB -u TiB	--size-unit MB --size-unit GB --size-unit TB --size-unit MiB --size-unit GiB --size-unit TiB	Display the total storage capacity of each volume and all other storage devices together, in a particular size unit, as specified by the parameter typed after the <code>-u</code> or <code>--size-unit</code> argument. The possible size units are: MB, GB, TB, MiB, GiB, or TiB.
-H		Display the storage volume and machine IDs in hexadecimal base.
-f <file name> -t csv -t xml	--file <file name> --out csv --out xml	Save the <code>xiv_devlist</code> information to a file. The name of the file must be typed after the <code>-f</code> or <code>--file</code> argument, and should be followed by the <code>-t</code> or <code>--out</code> argument that determines the file format: CSV or XML. For example: <pre>xiv_devlist -f file_name -t csv</pre>
-t csv -t xml	--out csv --out xml	Display the <code>xiv_devlist</code> information on the screen (without saving to a file), in either CSV or XML format (as specified). For the CVS output option, you can add the <code>--no-headers</code> option in the command line to generate the CSV output without column headers.
-h	--help	Display usage format help information that is specific to <code>xiv_devlist</code> .

Table below details the codes returned by the `xiv_devlist` command to specify the execution status.

Table 7. `xiv_devlist` exit codes

Code	Description
0	Successful execution
1	General error
2	Command executed with wrong arguments
4	Logging error

Setting Fibre Channel connectivity parameters

The `xiv_fc_admin` utility allows you to perform different administrative operations for FC-connected hosts and IBM storage systems.

Any operation performed with `xiv_fc_admin` requires a specific command argument, as detailed in Table 8 on page 74.

Important: All command arguments are case-sensitive.

Attention: If you are using AIX 6.1 TL3 or earlier, you must reboot the host after using **xiv_fc_admin -C**.

Table 8. Arguments for **xiv_fc_admin**

Short-form argument	Long-form argument	Use after xiv_fc_admin to:
-P	--print	Display the world-wide port number (WWPN) of the host bus adapters (HBAs) that are installed on this host.
-D -u <user name> -p <password> -H <hostname> -S <serial> -o <domain>	--define (followed by the sub-arguments)	<p>Define this host on a specific storage system, as detailed by the sub-arguments and parameters that you type after the -D or --define argument. The sub-arguments are:</p> <ul style="list-style-type: none"> • User name (-u) – User name for accessing the storage system on which the host should be defined. • Password (-p) – Password for accessing the storage system on which the host should be defined. • Hostname (-H) – Name of this host to be defined on the storage system. This sub-argument is optional. If not specified, the operating system's hostname is used by default. • Serial (-S) – Serial number of the specific storage system on which this host should be defined. Use the xiv_syslist command to display the serial numbers of all connected storage systems. • Domain (-o) – The storage domain or domains in which this host should be defined. This option is applicable only to XIV microcode version 11.5.0 or later, which supports the multi-tenancy feature. You can specify more than one domain by using a separating comma between the domain names. For example: <pre>xiv_iscsi_admin -D -o domain1, domain2</pre> <p>Important: If you are defining the host on an domain, you must provide credentials of a storage administrator that is attached to the domain (user name and password; see the -u and -p arguments above).</p>
-m	-m IGNORE -m DEFAULT	<p>Enter -m IGNORE to use a third-party multipath framework or enter -m DEFAULT to utilize the native multipath framework of the operating system. The -m attribute is used with the -R, -C, or -V flags.</p> <p>This argument must be typed on the command line when both the native multipath framework and a non-native multipath framework exist on the host.</p> <p>Note: This option is applicable only to Linux hosts.</p>
-R	--rescan	<p>Rescan and detect all FC-connected storage systems.</p> <p>Important: Always use xiv_fc_admin -R after new LUNs are added on the storage system (if FC connectivity is used).</p> <p>Optional: you can use the --clean argument with -R to disassociate unreachable devices from the host.</p>

Table 8. Arguments for `xiv_fc_admin` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_fc_admin</code> to:
-C	--configure	Check whether this host is properly configured for FC host attachment, and automatically perform any required configuration. The following is an example of the displayed output when using <code>xiv_fc_admin -C</code> on a Linux host: <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin-top: 10px;"> <pre>Configuring multipath - multipath.conf OK Configuring multipath service(s) OK</pre> </div>
-V	--verify	Verify the host configuration tasks.
-b	--build	Display the HAK build number.
-v	--version	Display the HAK version number. Note: You can use this argument with <code>xiv_attach</code> as well.
-i	--verbose	Display detailed information about the changes performed on the host by <code>xiv_fc_admin</code> , without carrying out the command. For more information, see “Displaying information about host configuration changes” on page 79.
-h	--help	Display usage format help information that is specific to <code>xiv_fc_admin</code> . Note: You can use this argument with <code>xiv_attach</code> as well.

Table below details the codes returned by the `xiv_fc_admin` command to specify the execution status.

Table 9. `xiv_fc_admin` exit codes

Code	Description
0	Successful execution
1	General error
2	Command executed with wrong arguments
3	Successful execution, but the host must be rebooted
4	Logging error
5	Verification stage failed

Setting iSCSI connectivity parameters

The `xiv_iscsi_admin` utility allows you to perform different administrative operations for iSCSI-connected hosts and IBM storage systems.

AIX

AIX users

Note: Because the IBM Storage Host Attachment Kit for AIX does not support iSCSI connectivity, refer to “Setting Fibre Channel connectivity parameters” on page 73.

HP-UX

HP-UX users

Note: Because the IBM Storage Host Attachment Kit for HP-UX does not support iSCSI connectivity, refer to “Setting Fibre Channel connectivity parameters” on page 73.

Any operation performed with **xiv_iscsi_admin** requires a specific command argument, as detailed in Table 10.

Important: All command arguments are case-sensitive.

Table 10. Arguments for **xiv_iscsi_admin**

Short-form argument	Long-form argument	Use after xiv_iscsi_admin to:
-c <IP address> [-n <CHAP name>] [-s <CHAP secret>]	--connect (followed by the IP address and CHAP sub-arguments if needed)	Connect to a new iSCSI target, specified by IP interface address that you type after the -c (small-cap c) or --connect argument. In addition, if CHAP authentication was defined for the iSCSI target (for more information, see “Using CHAP authentication for iSCSI connectivity” on page 21), use the following sub-arguments (do not use if CHAP authentication was not defined): <ul style="list-style-type: none"> • CHAP name (-n) – The CHAP authentication name. • CHAP secret (-s) – The CHAP authentication secret. Note: Using this argument with an existing connection will connect new targets (for example, when adding interface modules).
-d <IP address>	--disconnect	Disconnect an iSCSI target, specified by IP interface address that you type after the -d or --disconnect argument.
-P	--print	Display the iSCSI initiator name and alias name.
-D -u <user name> -p <password> -H <hostname> -S <serial> -o <domain>	--define	Define this host on a specific storage system, as detailed by the sub-arguments and parameters that you type after the -D or --define argument. The sub-arguments are: <ul style="list-style-type: none"> • User name (-u) – User name for accessing the storage system on which the host should be defined. • Password (-p) – Password for accessing the storage system on which the host should be defined. • Hostname (-H) – Name of this host to be defined on the storage system. This sub-argument is optional. If not specified, the operating system's hostname is used by default. • Serial (-S) – Serial number of the specific storage system on which this host should be defined. Use the xiv_syslist command to display the serial numbers of all connected storage systems. • Domain (-o) – The XIV domain or domains in which this host should be defined. This option is applicable only to XIV microcode version 11.5.0 or later, which supports the XIV multi-tenancy feature. You can specify more than one XIV domain by using a separating comma between the domain names. For example: <div style="border: 1px solid black; border-radius: 15px; padding: 5px; margin: 10px 0;"> xiv_iscsi_admin -D -o domain1, domain2 </div> Important: If you are defining the host on an XIV domain, you must provide credentials of an XIV administrator that is attached to the domain (user name and password; see the -u and -p arguments above).

Table 10. Arguments for `xiv_iscsi_admin` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_iscsi_admin</code> to:
<p>-R -m <DEFAULT>, <IGNORE>, <SWITCH></p>	<p>--rescan</p>	<p>Rescan and detect all iSCSI-connected storage systems. If the host-side acceleration (HSA) is enabled, use the multipath mode sub-argument (-m) to modify the rescan process as follows:</p> <ul style="list-style-type: none"> • -m SWITCH– Rescan is performed, taking into account the <i>dm-switch</i> mapping. • -m IGNORE– Rescan is performed, but the <i>dm-switch</i> mapping is not updated. <p>If the HSA is disabled, use -m DEFAULT, or the -R flag without any sub-argument to rescan and detect all iSCSI-connected storage systems. Important: Always use <code>xiv_iscsi_admin -R</code> after new LUNs are added on the storage system (if iSCSI connectivity is used).</p> <p>Optional: you can use the --clean argument with -R to disassociate unreachable devices from the host.</p>
<p>-C -m <DEFAULT>, <IGNORE>, <SWITCH></p>	<p>--configure</p>	<p>Check whether this host is properly configured for iSCSI host attachment, and automatically perform any required configuration. The configuration process can be modified using the multipath mode sub-argument (-m), as detailed below:</p> <ul style="list-style-type: none"> • -m DEFAULT– Multipath and iSCSI parameters are configured, <i>dm-switch</i> parameters are ignored. <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin: 5px 0;"> <pre>Configuring multipath - multipath.conf OK Configuring multipath service(s) OK Configuring iSCSI initiator name - initiatorname.iscsi OK Configuring iSCSI daemon - iscsid.conf OK Configuring iSCSI service OK</pre> </div> • -m IGNORE– iSCSI parameters are configured, multipath and <i>dm-switch</i> parameters are ignored. <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin: 5px 0;"> <pre>Configuring iSCSI initiator name - initiatorname.iscsi OK Configuring iSCSI daemon - iscsid.conf OK Configuring iSCSI service OK</pre> </div> • -m SWITCH– <i>dm-switch</i>, multipath and iSCSI parameters are configured. <div style="border: 1px solid black; border-radius: 10px; padding: 5px; margin: 5px 0;"> <pre>Configuring dm-switch kernel module installation OK Configuring dm-switch cron job OK Configuring dm-switch systemd service OK Configuring multipath - multipath.conf OK Configuring multipath service(s) OK Configuring iSCSI initiator name - initiatorname.iscsi OK Configuring iSCSI daemon - iscsid.conf OK Configuring iSCSI service OK</pre> </div> <p>Note: This option is applicable only to Linux hosts.</p>

Table 10. Arguments for `xiv_iscsi_admin` (continued)

Short-form argument	Long-form argument	Use after <code>xiv_iscsi_admin</code> to:																																
-V -m <DEFAULT> , <IGNORE> , <SWITCH>	--verify	Verify the host configuration tasks. The verification process can be modified, using the multipath mode sub-argument (-m), as detailed below: <ul style="list-style-type: none"> • -m DEFAULT– Multipath and iSCSI configuration tasks are verified, <i>dm-switch</i> tasks are ignored. <table border="1" style="margin-left: 20px;"> <tr><td>Verifying multipath - multipath.conf</td><td>OK</td></tr> <tr><td>Verifying multipath service(s)</td><td>OK</td></tr> <tr><td>Verifying iSCSI initiator name - initiatorname.iscsi</td><td>OK</td></tr> <tr><td>Verifying iSCSI daemon - iscsid.conf</td><td>OK</td></tr> <tr><td>Verifying iSCSI service</td><td>OK</td></tr> </table> • -m IGNORE– iSCSI configuration tasks are verified, multipath and <i>dm-switch</i> tasks are ignored. <table border="1" style="margin-left: 20px;"> <tr><td>Verifying iSCSI initiator name - initiatorname.iscsi</td><td>OK</td></tr> <tr><td>Verifying iSCSI daemon - iscsid.conf</td><td>OK</td></tr> <tr><td>Verifying iSCSI service</td><td>OK</td></tr> </table> • -m SWITCH– <i>dm-switch</i>, multipath and iSCSI configuration tasks are verified. <table border="1" style="margin-left: 20px;"> <tr><td>Verifying dm-switch kernel module installation</td><td>OK</td></tr> <tr><td>Verifying dm-switch cron job</td><td>OK</td></tr> <tr><td>Verifying dm-switch systemd service</td><td>OK</td></tr> <tr><td>Verifying multipath - multipath.conf</td><td>OK</td></tr> <tr><td>Verifying multipath service(s)</td><td>OK</td></tr> <tr><td>Verifying iSCSI initiator name - initiatorname.iscsi</td><td>OK</td></tr> <tr><td>Verifying iSCSI daemon - iscsid.conf</td><td>OK</td></tr> <tr><td>Verifying iSCSI service</td><td>OK</td></tr> </table> <p>Note: This option is applicable only to Linux hosts.</p>	Verifying multipath - multipath.conf	OK	Verifying multipath service(s)	OK	Verifying iSCSI initiator name - initiatorname.iscsi	OK	Verifying iSCSI daemon - iscsid.conf	OK	Verifying iSCSI service	OK	Verifying iSCSI initiator name - initiatorname.iscsi	OK	Verifying iSCSI daemon - iscsid.conf	OK	Verifying iSCSI service	OK	Verifying dm-switch kernel module installation	OK	Verifying dm-switch cron job	OK	Verifying dm-switch systemd service	OK	Verifying multipath - multipath.conf	OK	Verifying multipath service(s)	OK	Verifying iSCSI initiator name - initiatorname.iscsi	OK	Verifying iSCSI daemon - iscsid.conf	OK	Verifying iSCSI service	OK
Verifying multipath - multipath.conf	OK																																	
Verifying multipath service(s)	OK																																	
Verifying iSCSI initiator name - initiatorname.iscsi	OK																																	
Verifying iSCSI daemon - iscsid.conf	OK																																	
Verifying iSCSI service	OK																																	
Verifying iSCSI initiator name - initiatorname.iscsi	OK																																	
Verifying iSCSI daemon - iscsid.conf	OK																																	
Verifying iSCSI service	OK																																	
Verifying dm-switch kernel module installation	OK																																	
Verifying dm-switch cron job	OK																																	
Verifying dm-switch systemd service	OK																																	
Verifying multipath - multipath.conf	OK																																	
Verifying multipath service(s)	OK																																	
Verifying iSCSI initiator name - initiatorname.iscsi	OK																																	
Verifying iSCSI daemon - iscsid.conf	OK																																	
Verifying iSCSI service	OK																																	
-U	--update	Update routing table on the device-mapper multipath switch (<i>dm-switch</i>). This command is available for RHEL 7.x with the host-side acceleration (HSA) enabled.																																
-b	--build	Display the HAK build number.																																
-v	--version	Display the HAK version number. Note: You can use this argument with <code>xiv_attach</code> as well.																																
-i	--verbose	Display detailed information about the changes performed on the host by <code>xiv_iscsi_admin</code> , without carrying out the command. For more information, see “Displaying information about host configuration changes” on page 79.																																
-h	--help	Display usage format help information that is specific to <code>xiv_iscsi_admin</code> . Note: You can use this argument with <code>xiv_attach</code> as well.																																

Table below details the codes returned by the `xiv_iscsi_admin` command to specify the execution status.

Table 11. `xiv_iscsi_admin` exit codes

Code	Description
0	Successful execution
1	General error

Table 11. `xiv_iscsi_admin` exit codes (continued)

Code	Description
2	Command executed with wrong arguments
3	Successful execution, but the host must be rebooted
4	Logging error
5	Verification stage failed

Displaying information about host configuration changes

This section describes the use and output of the `-i` or `--verbose` command argument, when used with HAK commands on Linux or Windows Server hosts.

Linux

Note: This section is relevant to Linux and Windows Server users only.

Windows Server

If you want to display a detailed description of the configuration changes that are to be performed when running the `xiv_attach`, `xiv_fc_admin`, and `xiv_iscsi_admin` commands, use the `-i` or `--verbose` argument after the command, as described in the following sections:

- “Displaying information for `xiv_fc_admin` or `xiv_iscsi_admin`”
- “Displaying information when working with `xiv_attach`” on page 80

For a summary of all possible configuration changes on the host, refer to Chapter 9, “Summary of changes applied on the host,” on page 97.

Displaying information for `xiv_fc_admin` or `xiv_iscsi_admin`

When using the `-i` or `--verbose` argument after `xiv_fc_admin` or `xiv_iscsi_admin`, the details of the host configuration actions that these commands perform are displayed, but these actions are not carried out.

You can review the displayed information and then decide whether you want to use the `xiv_fc_admin --configure` or `xiv_iscsi_admin --configure` commands on the host (for more information about the `--configure` argument, see Table 8 on page 74 and Table 10 on page 76).

The following example shows the displayed information for `xiv_fc_admin` on a Linux host:

```
xiv_fc_admin --verbose

multipath - multipath.conf
=====
The multipath configuration file (/etc/multipath.conf) is updated as follows:

- The IBM storage device is added/updated in the devices section, along with
its recommended parameters.
- Any blacklisting of the IBM storage device is removed.

Before applying these changes, a backup copy of the original multipath.conf
file is saved in <HAK_dir>/backup.

On RHEL 6.3 or later, the device-mapper-multipath package already includes the
required IBM storage system values, so the multipath.conf
file is not modified.

multipath service(s)
=====
The multipath service ("multipathd") is enabled on run levels 3, 4, and 5, and
started automatically if notcurrently running.

On SLES, the 'boot.multipath' service is enabled on the 'boot' run level.
```

For more information about the configuration operations on the host per HAK command and operating system, see Chapter 9, “Summary of changes applied on the host,” on page 97.

Displaying information when working with `xiv_attach`

When using the `-i` or `--verbose` argument after `xiv_attach`, information about pending host configuration changes is displayed, followed by a confirmation prompt.

The configuration changes are carried out only after your confirmation.

The following example shows the displayed information for `xiv_attach` on a Windows Server 2008/2012 host that uses iSCSI connectivity:

```
xiv_attach --verbose

Previous HAK versions
=====
All previous versions of the IBM Storage Host Attachment Kit are removed from the
Windows registry, and all previous versions of Python are removed from 'pythonpath'.

LUN0 device driver
=====
The LUN0 device driver is installed to enable display of IBM storage systems in
the list of recognized devices.

Disk timeout setting
=====
The value of the 'TimeOutValue' key in the Windows Server registry is set to 60.
The key is located on the following path:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk

iSCSI service
=====
The 'Microsoft Software iSCSI Initiator' service is set to start automatically upon
boot, and set to be started manually upon the first use.

Persistent reservation key
=====
A persistent reservation key is set so that a Round-Robin MPIO policy could be used
when using the IBM Storage Enabler for Windows Failover Clustering.
```

For more information about the configuration operations on the host per HAK command and operating system, see Chapter 9, “Summary of changes applied on the host,” on page 97.

Chapter 7. Diagnostics

The IBM Storage Host Attachment Kit (HAK) includes two utilities that allow you to perform comprehensive diagnostics of the host, as explained in the following sections.

- “Using the host diagnostics utility”
- “Using the Host Profiler utility” on page 85

Using the host diagnostics utility

Whenever you want to perform complete diagnostics of your host and its connectivity with the IBM storage system, run the **xiv_diag** utility.

The gathered information is saved to a file that you can review locally or send to customer support whenever further assistance is needed.

Note: The gathered information should be interpreted and analyzed by an IBM professional. The generated file does not include any analysis of the gathered information.

The following is an example of the displayed output when using **xiv_diag**:

```
xiv_diag

Welcome to the IBM diagnostics tool, version 2.6.0.
This tool will gather essential support information from this host and save the information
to a file.

Please specify the directory into which the xiv_diag file should be saved [default: /tmp]:
```

When prompted to specify a directory, either press Enter to generate the report file in the default folder (/tmp in Linux as shown in the example above), or enter a different path.

Alternatively, if you want to define the directory in advance and avoid the directory prompt, use the **-p** or **--path** argument right after the **xiv_diag** command, and then specify the folder path. This eliminates the need for entering the folder path in a second step. For example:

```
xiv_diag -p /var/log/
```

To list the gathered information on the screen without collecting it or saving it to a file (in cases when the file generation takes a long time and the file is not needed), type **-L** after **xiv_diag** on the command line:

```
xiv_diag -L
```

When the information gathering starts, you can view its progress as shown in the following Linux-specific example (the output varies depending on the host operating system and HAK version):

```

Creating archive xiv_diag-results_2014-1-8_11-34-2
INFO: Gathering HAK version... DONE
INFO: Gathering lsb_release -a... DONE
INFO: Gathering xiv_devlist... DONE
INFO: Gathering xiv_fc_admin -L... DONE
INFO: Gathering xiv_fc_admin -V... DONE
INFO: Gathering xiv_fc_admin -P... DONE
INFO: Gathering xiv_iscsi_admin -L... DONE
INFO: Gathering xiv_iscsi_admin -V... DONE
INFO: Gathering xiv_iscsi_admin -P... DONE
INFO: Gathering multipath... DONE
INFO: Gathering ls scsi devices... DONE
INFO: Gathering fdisk -l... DONE
INFO: Gathering modinfo qla2xxx... DONE
INFO: Gathering modinfo qla4xxx... DONE
INFO: Gathering modinfo bfa... DONE
INFO: Gathering modinfo lpfc... DONE
INFO: Gathering df... DONE
INFO: Gathering uname -a... DONE
INFO: Gathering dmidecode... DONE
INFO: Gathering rpm -qa... DONE
INFO: Gathering ifconfig -a... DONE
INFO: Gathering inquiry.py... DONE
INFO: Gathering List disks in /dev directory... DONE
INFO: Gathering /etc/fstab... DONE
INFO: Gathering /var/lib/multipath/bindings... DONE
INFO: Gathering /etc/multipath/bindings... DONE
INFO: Gathering /etc/initiatorname.iscsi... DONE
INFO: Gathering /etc/iscsi/initiatorname.iscsi... DONE
INFO: Gathering /etc/iscsi.conf... DONE
INFO: Gathering /etc/iscsid.conf... DONE
INFO: Gathering /etc/iscsi/iscsid.conf... DONE
INFO: Gathering multipath -t... DONE
INFO: Gathering multipath.conf... DONE
INFO: Gathering mtab... DONE
INFO: Gathering /etc/issue... DONE
INFO: Gathering /etc/redhat-release... DONE
INFO: Gathering build-revision file... DONE
INFO: Gathering /sys directory... DONE
INFO: Gathering /etc/udev directory... DONE
INFO: Gathering /proc directory... DONE
INFO: Gathering /etc/iscsi directory... DONE
INFO: Gathering /var/lib/iscsi directory... DONE
INFO: Gathering /var/log/messages logs... DONE
INFO: Gathering /var/log/boot logs... DONE
INFO: Gathering host_attach logs... DONE
INFO: Gathering xiv_devlist logs... DONE
INFO: Closing xiv_diag archive file... DONE
Deleting temporary directory... DONE
INFO: Information gathering has been completed.
INFO: You can now send /tmp/xiv_diag-results_2014-1-8_11-34-2.tar.gz to
IBM Support for further review.

```

Note: Gathered information items are marked as 'DONE'. Some items may be skipped and marked as 'SKIPPED' due to technical incompatibilities with the host operating system, or due to other technical reasons.

Table below details the codes returned by the **xiv_diag** command to specify the execution status.

Table 12. xiv_diag exit codes

Code	Description
0	Successful execution
1	General error
2	Command executed with wrong arguments

Using the Host Profiler utility

The Host Profiler utility (**xiv_host_profiler**) collects comprehensive configuration information from the host on which it is run, and creates a profile based on the collected information.

By default, the host profile is saved to a file and uploaded to supported IBM storage systems. The profile can then be shared with different IBM internal technical support systems.

Note: The Host Profiler utility can run on any host, but can upload information only to XIV storage systems of microcode version 11.2 or later.

The created host profile includes different information parameters, as detailed in the following table.

Table 13. Host profile information parameters

Information parameter	Description
Connected system microcode version(s)	Microcode versions of connected IBM storage systems (single system or multiple systems) to which the host is connected.
Operating system name	Name of the operating system that runs on the host.
Operating system version	Version of the operating system that runs on the host
Operating system patch level	The installed patch level or service pack of the operating system that runs on the host.
FC connectivity used	Indicates whether the host is connected to a storage system over Fibre Channel.
iSCSI connectivity used	Indicates whether the host is connected to a storage system over iSCSI.
Multipath framework(s)	Types and versions of the multipath frameworks that are installed and configured on the host.
HBA information	Information about the host bus adapters (HBAs) that are installed on the host. Includes each HBA model name, driver version, and firmware version.
Host CPU architecture	CPU architecture of the host.
Is the host connected to two or more modules on each storage system	Indicates whether there are connectivity paths to at least two separate modules of each storage system to which the host is connected.
Hotfixes installed on the host	A list of vendor-issued hotfixes that are installed on the host. Applicable only to Microsoft Windows Server.
IBM Storage Host Attachment Kit version	Version and build number of the IBM Storage Host Attachment Kit with which the host profile was created.
IBM Storage Enabler for Windows Failover Clustering version	Version of the IBM Storage Enabler for Windows Failover Clustering ("MSCS Agent"), if installed on the host. Applicable only to Microsoft Windows Server.
Storport driver version	Version of the storport.sys driver file that is installed on the host. Applicable only to Microsoft Windows Server.
Host profile created at	Date and time at which the host profile was created.
Host unique ID	Hardware-based unique identifier of the host. Helps remote support systems make a distinction between identical host profiles, or monitor how the profile of a certain host changes over time. The host ID remains the same even after OS upgrades or reinstallation.

Table 13. Host profile information parameters (continued)

Information parameter	Description
Was the host profile committed to all IBM storage systems	Indicates whether this specific host profile was uploaded to all the connected storage systems.
Serialized host profile	A serialized representation of the host profile, written to the storage system. When using --analyze with --file , the Host Profiler utility reads this value and analyzes it, allowing analysis of profiles that were created on different hosts.

In addition to sharing the collected information with other systems, the Host Profiler utility includes a detailed analysis function that can help you determine whether the host is properly configured for operation with storage systems, and learn what actions might help optimize the configuration or resolve detected issues.

Note: The analysis operation is performed against a bundled preset rule file that includes a list of tested and approved IBM configurations. This rule file is updated from time to time. You can obtain the latest version of it from the IBM Fix Central (www.ibm.com/support/fixcentral).

To streamline the remote support process for the host, you can set the utility to run on a weekly basis by using the scheduled task function. Such scheduling is recommended as a best practice. For more information, refer to Table 14.

Any operation performed with **xiv_host_profiler** requires a specific command argument, as detailed in Table 14.

Important: All command arguments are case-sensitive.

Table 14. Arguments for **xiv_host_profiler**

Argument	Use after xiv_host_profiler to:	Additional options
--create	<p>Create a host profile and upload it to the connected storage system (or systems). In addition, an XHP file containing the profile information is created in the operating system default log directory.</p> <p>For example:</p> <pre>xiv_host_profiler --create</pre> <p>Host Profile written to /var/log/hostname.xhp</p>	<ul style="list-style-type: none"> • Type --local after --create to create only the local XHP file, without uploading the profile to the storage system. • Type --stored after --create to upload the profile to storage systems without creating a local file. • Add -f <file name> or --file <file name> to the command line in order to specify the directory path and name of the XHP file that should store the created host profile.

Table 14. Arguments for `xiv_host_profiler` (continued)

Argument	Use after <code>xiv_host_profiler</code> to:	Additional options
<p>--analyze</p>	<p>Perform a comprehensive analysis of the host configuration in order to determine whether the storage connectivity components are properly installed and supported by IBM.</p> <p>The analysis is performed against a rule file (<code>xhop_audit.dat</code>) bundled with the IBM Storage Host Attachment Kit. The date of the last rule file update appears at the end of the analysis output.</p> <p>The returned analysis information depends on the XIV microcode versions that have been detected as connected to the host.</p> <p>For example:</p> <pre>xiv_host_profiler --analyze Host configuration analysis completed. Results for microcode version 11.2.X (detected as currently in use): The host is not configured for optimal performance with IBM XIV storage systems software version 11.2.X The following problems were detected: 1. Problem 1 2. Problem 2 Configuration notes: 1. Note 1 2. Note 2 Results for microcode version 11.4.X (detected as currently in use): The host is configured for optimal performance with IBM XIV storage systems software version 11.4.X Configuration notes: 1. Note 1 2. Note 2 The above analysis is based on a rule file updated on 15 November 2013. The latest rules file can be downloaded from https://ibm.biz/Bdx8Qc.</pre>	<ul style="list-style-type: none"> • Type --xiv-microcode <microcode version> after --analyze to perform analysis against a specific XIV microcode version (whether a detected version or non-detected version). For example: <pre>xiv_host_profiler --analyze --xiv-microcode 11.2.0</pre> • Type --rules <file name> after --analyze to perform analysis against a rule file that is different than the rule file bundled with the IBM Storage Host Attachment Kit. For example: <pre>xiv_host_profiler --analyze --rules ruledir/rules.dat</pre> • Add -f <file name> or --file <file name> to the command line in order to analyze a previously saved host profile (of the same host or any other host).
<p>--create-scheduled-task</p>	<p>Create a scheduled task (cronjob) that updates the host profile each week on Sunday, at 4:00 AM. Note: Creating this scheduled task is recommended as a best practice.</p>	<p>You can also create the same scheduled task when running the xiv_attach command, as explained in “Stage 1: Using the CLI attachment utility (<code>xiv_attach</code>)” on page 37.</p>
<p>--remove-scheduled-task</p>	<p>Remove the previously set scheduled task that updates the host profile each week. This option removes only the task that was created by --create-scheduled-task. Note: Tasks that were set manually (not by using the --create-scheduled-task command) are not removed.</p>	<p>N/A</p>
<p>--version (or -v)</p>	<p>Display the HAK version number.</p>	<p>N/A</p>

Table 14. Arguments for `xiv_host_profiler` (continued)

Argument	Use after <code>xiv_host_profiler</code> to:	Additional options
<code>--help</code> (or <code>-h</code>)	Display usage format help information that is specific to <code>xiv_host_profiler</code> .	N/A

Table below details the codes returned by the `xiv_host_profiler` command to specify the execution status.

Table 15. `xiv_host_profiler` exit codes

Code	Description
0	Successful execution/Host profile analysis completed with optimal results
1	General error
2	Operating system not supported
3	Failed to upload the host profile to IBM storage system
4	Failed writing the host profile to file
5	Host profile analysis completed with non-optimal results
6	Host profile analysis completed with inconclusive results

Chapter 8. Best practices

There are a few best practices that you can apply when using the IBM Storage Host Attachment Kit, as described in the following sections.

- “Recommended operating system settings”
- “Monitoring host connectivity”
- “Using the rescan (-R) option after new volumes are mapped” on page 92
- “Scheduling the Host Profiler utility for weekly operation” on page 92
- “Using the host-side acceleration (HSA)” on page 93

Recommended operating system settings

Each operating system and CPU architecture has different settings that work best for host connectivity and performance with the IBM storage system.

If you are a host administrator, consult with your storage administrator for setting up the best-suited host configuration in your specific production environment.

For detailed information about recommended settings refer to *IBM XIV Storage System – Host Attachment and Interoperability* (www.redbooks.ibm.com/abstracts/sg247904.html), available on the IBM Storage Redbooks website (www.redbooks.ibm.com/redbooks.nsf/portals/storage).

Monitoring host connectivity

Constant host connectivity is crucial for hosts that work with the IBM storage system.

To ensure that your host is properly connected to the storage system, you can monitor the connectivity of attached hosts in real time from either the XIV GUI or XCLI. In addition, you can check host connectivity from the IBM Hyper-Scale Manager.

Monitoring from the XIV GUI

This section describes how to use the XIV GUI for monitoring.

Procedure

1. Log in to the IBM XIV GUI management application and zoom into the XIV storage system to which the host is connected.
2. Move the mouse over the **Hosts and Clusters** icon and then click **Hosts Connectivity**.



Figure 51. XIV GUI - Hosts Connectivity

The Hosts Connectivity table is displayed.

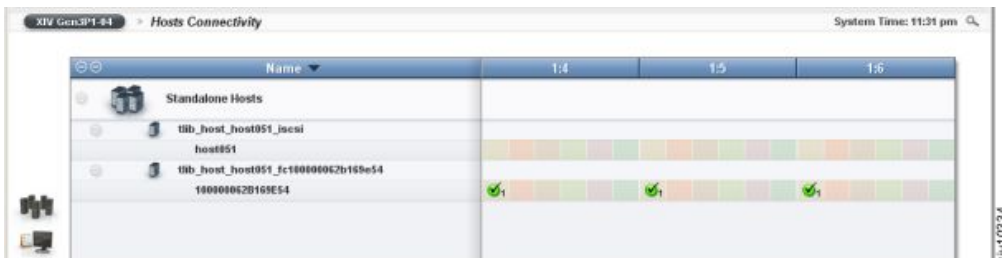


Figure 52. XIV GUI - Hosts Connectivity table

Monitoring from the CLI

If you want to monitor the host connectivity from the CLI rather than from the GUI, use the following commands.

- List FC and iSCSI connectivity of a host:

```
host_connectivity_list [ host=HostName | fc_host_port=WPN ]
[module=ComponentId | fcport=ComponentId ]
```

- List the status and configuration of the storage system FC ports:

```
fc_port_list [ module=ComponentId | fcport=ComponentId ]
```

For more detailed information about CLI commands, refer to:

- IBM XIV XCLI Reference Guide, available in the IBM XIV Storage System documentation(www.ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html).
- IBM Spectrum Accelerate Command-Line Interface (CLI) Reference Guide, available in the IBM Spectrum Accelerate documentation(http://www.ibm.com/support/knowledgecenter/STZSWD/landing/IBM_Spectrum_Accelerate_welcome_page.html).

Monitoring from the IBM Hyper-Scale Manager

This section describes how to use the XIV GUI for monitoring.

Procedure

1. Log in to the IBM Hyper-Scale Manager application and zoom into the storage system to which the host is connected.
2. Click the **Hosts and Clusters** icon on the sidebar and then click **Hosts Connectivity**.
The Hosts Connectivity table is displayed.




Figure 53. Hosts & Clusters Views menu

HOST Click here to adjust filter						
Host ^	Type	System	Domains	Cluster	Volumes	Snapshots
111	Default	XIV_Gen3P1-01a	/Global Space/		2	0
1111	Default	XIV_Gen3P1-01a	/Global Space/		2	0
111111111111111111...	Default	XIV_Gen4D-27a	/Global Space/		0	0
12	HPUX	XIV_Gen3P1-01a	/Global Space/		0	1
2	HPUX	XIV_Gen4D-27a	/Global Space/		0	0
2345	Default	XIV_MTGenP-01	/Global Space/		2	0
27a	Default	XIV_Gen4D-27a	/Global Space/		0	0
5	Default	XIV_Gen4D-27a	/Global Space/	shaked	0	0
aaa	Default	XIV_6012019	/Global Space/		0	0
asd	Default	XIV_Gen4D-27a	/Global Space/		1	0
AShaked	Default	XIV_Gen4D-27a	2	AShaked	0	0
atmp	Default	XIV_Gen4D-27a	/Global Space/	ErezCluster2	0	0
b	Default	XIV_Gen3P1-01a	2		1	0
dan	Default	XIV_Gen3P1-01a	/Global Space/		0	0
danny	Default	2	3	shaked	1	0
danny1	Default	XIV_Gen3P1-01a	/Global Space/		0	0

Host 111 Statistics: IOPS (total) 0 iops ↑ Latency (avg) 0.00 ms ↑ Bandwidth (avg) 0.00 mb/s ↑

Figure 54. Hosts Connectivity table

Using the rescan (-R) option after new volumes are mapped

In a typical environment, new storage volumes may be created on a regular basis.

As a best practice and to avoid missing new volumes, always use the `xiv_iscsi_admin -R` or `xiv_fc_admin -R` command (depending on the connectivity type) to rescan for volumes after new volumes are mapped to the host (for more information, see “Stage 3: Mapping storage volumes to the host” on page 49, “Setting Fibre Channel connectivity parameters” on page 73, and “Setting iSCSI connectivity parameters” on page 75).

Scheduling the Host Profiler utility for weekly operation

To streamline the remote support process for the host, you can set the utility to run on a weekly basis by using the scheduled task function. Such scheduling is recommended as a best practice.

Use the `--create-scheduled-task` option with the Host Profiler utility, as explained in “Using the Host Profiler utility” on page 85, or schedule the task during the attachment procedure, as explained in “Stage 1: Using the CLI attachment utility (`xiv_attach`)” on page 37.

Using the host-side acceleration (HSA)

This section describes how to use the host-side acceleration (HSA), available for Spectrum Accelerate storage systems.

The HSA improves iSCSI communication between Spectrum Accelerate and a host computer. The improvement is achieved by efficient mapping of I/O operations to modules that store the required data. Instead of sending data across all paths and letting the storage arrays decide on correct forwarding, the host sends data directly to the correct module. The use of HSA dramatically reduces Ethernet traffic and decreases iSCSI latency. The HSA maintains a device-mapper switch (*dm-switch*) that relies on the host's multipath routing table for making routing decisions. The required routing data is retrieved from the volume distribution table on Spectrum Accelerate. This enables the host to learn which module holds the primary partition, serving the I/O, as illustrated in the figure below.

The HSA is supported on a host running RHEL 7.x and above, and connected to the Spectrum Accelerate storage system.




Figure 55. A host connected to IBM Spectrum Accelerate

The HSA can be enabled using one of the following methods:

- Entering the `xiv_attach` command to invoke the interactive CLI wizard. See “Stage 1: Using the CLI attachment utility (`xiv_attach`)” on page 37 for details.
- Entering the `xiv_iscsi_admin` command with `-C -m SWITCH` sub-argument. See “Setting iSCSI connectivity parameters” on page 75 for details.

In addition, the IBM Storage Host Attachment Kit runs a Linux cron job every 60 minutes to execute the `xiv_iscsi_admin -U` command for updating the `dm-switch` routing table. The update period can be adjusted manually by editing the configuration file.

If you have active Linux multipath devices, you can disengage the file system from them, and move it to the HSA multipath device. This procedure is explained in “Moving file system to the HSA `dm-switch`”

The HSA functionality can be disabled in the Linux environment, as detailed in “Disabling the HSA.”

Moving file system to the HSA `dm-switch`

This section describes how to move the existing file system from standard Linux mutlipath devices to the HSA `dm-switch`.

About this task

When the host is configured to use the HSA, a new `dm-switch` device is created for each volume. All block devices managed by the `dm-switch` share the same WWID, similar to standard Linux multipath structure. After enabling the HSA, you must re-mount the file system on the HSA multipath device, as explained below.

Note: Do not mount file systems that use `fdisk` partitioning on the HSA multipath devices. A data loss may occur, when the file system is moved to the standard Linux mutlipath devices.

Procedure

1. Stop all applications that use the `/dev/mapper/mpathx` multipath device.
2. Unmount `/dev/mapper/mpathx`.
3. Run the `multipath -f mpathx` command to remove the `/dev/mapper/mpathx` multipath device. You can also use the `multipath -f` command to clear all multipath devices.
4. Run the `xiv_iscsi_admin -U` command to create HSA multipath devices.
5. If you use the Linux Volume Manager (LVM), run the `lvmdiskscan` command to register the HSA multipath devices with the LVM.
6. Run the `mount -a` command to complete the mount procedure.
7. Restart all previously stopped applications.

Disabling the HSA

This section describes how to disable the host-side acceleration (HSA).

About this task

If no longer needed, you can disable the HSA, using the following procedure on Linux.

Procedure

1. Disable the HSA cron job.
`rm -f /etc/cron.d/xiv_switch_update`
2. Disable the HSA `systemd` service.

- ```
systemctl disable xiv-hsa.service
```
3. Remove the required WWIDs from blacklist in the `/etc/multipath.conf` file.
  4. If needed, remove any file system or LVM objects defined on the HSA devices manually, as detailed in the sub-steps below. You can also run a script to remove the HSA devices automatically, as detailed in step 5 Alternatively, you can move the file system to other multipath devices, as explained in “Moving file system from the HSA *dm-switch*.”
    - a. Remove all HSA switch devices: `#dmsetup ls --target switch`.
    - b. For each device, run: `#dmsetup remove <device_name>`.
    - c. Remove all HSA multipath devices (HSA multipath devices have the `ACC_` prefix): `#dmsetup ls --target multipath`.
  5. Run this script to remove the HSA devices automatically: `#python /opt/xiv/host_attach/lib/hak/host_attach/platform/linux/hsa/tools/clean.pyc`.
  6. Delete all files with the `md5` postfix from the `/tmp` directory.

### Moving file system from the HSA *dm-switch*

This section describes how to move the existing file system from the HSA *dm-switch*, after the HSA has been disabled.

#### About this task

After disabling the HSA, you must re-mount the file system on standard Linux mutlipath devices, as explained below. During this procedure, the LUNs used by the HSA *dm-switch* must be removed from the WWID blacklist section of the `multipath.conf` file to allow creation of the Linux mutlipath devices.

#### Procedure

1. Stop all applications that use the `/dev/mapper/ACC_xxx_yy` multipath device.
2. Unmount `/dev/mapper/ACC_xxx_yy` device.
3. If you use the Linux Volume Manager (LVM), run the `vgchange -an <vol_group_name>` command to deactivate the volume group.
4. Clear all existing HSA multipath devices. IBM recommends using this Python script to remove the HSA multipath devices: `#python /opt/xiv/host_attach/lib/hak/host_attach/platform/linux/hsa/tools/clean.pyc`.
5. Remove the WWID of the `/dev/mapper/ACC_xxx_yy` multipath device from the blacklist section of the `multipath.conf` file.
6. Restart the `multipathd` daemon.
7. Run the `multipath` command to create Linux multipath devices.
8. If you use the Linux Volume Manager (LVM), run the `lvm diskscan` command to register the new multipath devices with the LVM. If you do not use the LVM, add a relevant mountpoint, using newly created multipath devices in the `/etc/fstab` file.
9. Run the `mount -a` command to complete the mount procedure.
10. Restart all previously stopped applications.



---

## Chapter 9. Summary of changes applied on the host

Different changes are applied on the host when you perform different Host Attachment Kit operations.

This includes:

- Changes applied on a Linux host
- Changes applied on a Windows Server host

---

### Changes applied on a Linux host

This section summarizes the changes that can be applied to a Linux host.

**Linux**

**Note:** This section and the following subsections apply to Linux users only.

Depending on the Host Attachment Kit operations that you perform on the Linux host, the following changes may be applied on it:

- Multipath configuration
- Multipath service(s)
- iSCSI daemon configuration
- iSCSI initiator name
- iSCSI service

### Multipath configuration

**Note:** This subsection applies to the `xiv_fc_admin`, `xiv_iscsi_admin`, and `xiv_attach` commands when used on a Linux host.

The multipath configuration file (`/etc/multipath.conf`) is updated as follows:

- The storage device is added/updated in the devices section, along with its recommended settings and parameters:

**path\_checker tur**

The SCSI Test Unit Ready CDB is called for checking if path is up.

**no\_path\_retry queue**

Perform queue if all paths are down.

**rr\_min\_io 15**

Minimum I/Os before performing round-robin.

**path\_grouping\_policy multibus**

All paths should be used equally.

**path\_selector "round-robin 0" or "service-time 0"**

For RHEL 7.x, default service time is set to 0. For the previous RHEL versions, default round robin value is 0.

- Any blacklisting of the storage device is removed.

---

**Note:** When the host-side acceleration is enabled, the Spectrum Accelerate block devices are blacklisted. This enables the HSA module to make multipath decisions, instead of the *multipathd* service.

---

Before applying these multipath configuration changes, a backup copy of the original `multipath.conf` file is saved in `<HAK_dir>/backup`. If no `multipath.conf` file is present, it is automatically created (using `/sbin/mpathconf --enable`).

---

**Note:** On RHEL 6.3 or later and SLES 11.4 or later, **device-mapper-multipath** (version 0.4.9, release 56 or later) already includes the storage system defaults. If this version is detected, the `multipath.conf` file is not modified, unless the storage volumes are blacklisted in the file and need to be cleared. In such a case, the only change in `multipath.conf` is the removal of storage volumes from the black list.

---

## Multipath service(s)

---

**Note:** This subsection applies to the `xiv_fc_admin`, `xiv_iscsi_admin`, and `xiv_attach` commands when used on a Linux host.

---

The multipath service ("**multipathd**") is enabled on run levels 3, 4, and 5, and started automatically if not currently running.

On SLES, the `boot.multipath` service is enabled on the 'boot' run level.

## iSCSI daemon configuration

---

**Note:** This subsection applies to the `xiv_iscsi_admin` and `xiv_attach` commands when used on a Linux host and over iSCSI.

---

The iSCSI daemon configuration file (`/etc/iscsid.conf` or `/etc/iscsi/iscsid.conf`) is updated with the default values. The full list of values can be found in: `<HAK_dir>/etc/linux/iscsi.conf`

A backup copy of the original `iscsid.conf` is saved in `<HAK_dir>/backup`.

## iSCSI initiator name

---

**Note:** This subsection applies to the `xiv_iscsi_admin` and `xiv_attach` commands when used on a Linux host and over iSCSI.

---

A random IQN is generated and placed in: `/etc/[iscsi/]initiatorname.iscsi`

If the `initiatorname.iscsi` file already exists and contains an IQN, it is not modified. The random IQN is generated by the native Linux `iscsi-name` tool.

## iSCSI service

---

**Note:** This subsection applies to the `xiv_iscsi_admin` and `xiv_attach` commands when used on a Linux host and over iSCSI.

---



The iSCSI service is enabled on run levels 3, 4, and 5, and started automatically if not currently running. The service name is **iscsi** on RHEL, and **open-iscsi** on SLES.

---

## Changes applied on a Windows Server host

This section summarizes the changes that can be applied to a Windows server host.

### Windows Server

**Note:** This section and the following subsections apply to Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.

Depending on the Host Attachment Kit operations that you perform on the Windows Server host, the following changes may be applied on it:

- Removal of previous HAK versions
- LUN0 device driver
- Disk timeout setting
- iSCSI service
- XIV MPIO Load Balancing service
- XIV MPIO Load Balancing agent
- Built-in MPIO feature

### Removal of previous HAK versions

All previous versions of the IBM Storage Host Attachment Kit are removed from the Windows registry, and all previous versions of Python are removed from pythonpath.

### LUN0 device driver

The LUN0 device driver is installed to enable display of IBM storage systems in the list of recognized devices.

### Disk timeout setting

The value of the *TimeoutValue* key in the Windows Server registry is set to 60. The key is located on the following path:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
```

### iSCSI service

The **Microsoft Software iSCSI Initiator** service is set to start automatically upon boot, and set to be started manually upon the first use.

### IBM MPIO Load Balancing service

Any instance of the MPIO load balancing service is traced, stopped, and removed – as part of removing the deprecated IBM MPIO load balancing.

### IBM MPIO Load Balancing agent

Any instance of the MPIO load balancing agent is traced, stopped, and removed – as part of removing the deprecated IBM MPIO load balancing.

## **Built-in MPIO feature**

The built-in Microsoft Windows MPIO feature is installed and then enabled for IBM storage devices.

---

## Chapter 10. IBM storage systems and VMware ESXi hosts

The following reference information is relevant only to VMware hosts and not related to the IBM Storage Host Attachment Kit solution.

---

**Attention:** This chapter provides reference information about how to physically connect VMware ESXi hosts to the IBM storage system. However, **the Host Attachment Kit software solution is not available for VMware ESXi hosts, and cannot be used** with these hosts.

---

For more detailed information about IBM storage host software solutions for VMware platforms, refer to:

- IBM XIV Storage System on IBM Knowledge Center ([ibm.com/support/knowledgecenter/STJTAG](http://ibm.com/support/knowledgecenter/STJTAG))
- IBM Spectrum Accelerate on IBM Knowledge Center ([ibm.com/support/knowledgecenter/STZSWD](http://ibm.com/support/knowledgecenter/STZSWD))
- IBM FlashSystem A9000 on IBM Knowledge Center ([ibm.com/support/knowledgecenter/STJKMM](http://ibm.com/support/knowledgecenter/STJKMM))
- IBM FlashSystem A9000R on IBM Knowledge Center

---

### Connectivity types for ESXi hosts

The following section details the connectivity types supported ESXi platforms.

VMware ESXi 5.x and ESXi 6.x server platforms support Fibre Channel (FC) and iSCSI connectivity. For the latest support information and compatibility matrix, refer to the IBM System Storage Interoperation Center ([www.ibm.com/systems/support/storage/config/ssic](http://www.ibm.com/systems/support/storage/config/ssic)). Additional compatibility information is available in the VMware compatibility guide(<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san>).

For best performance, install the latest firmware and drivers for the HBAs that are in use. The HBA vendor should provide the latest firmware and drivers.

---

### Fibre Channel SAN zoning with ESXi hosts

SAN Fibre Channel zoning defines devices within a zone, so that it isolates devices inside a zone from devices outside the zone.

FC zoning is generally a good practice and a requirement when connecting ESXi servers with multiple paths to a storage system.

SAN zoning should be set as a Single Initiator Zoning layout. Set a single Initiator (ESXi HBA port) per SAN zone. Figure 56 on page 102 shows two clustered ESXi servers are attached to a storage system over multiple FC paths.

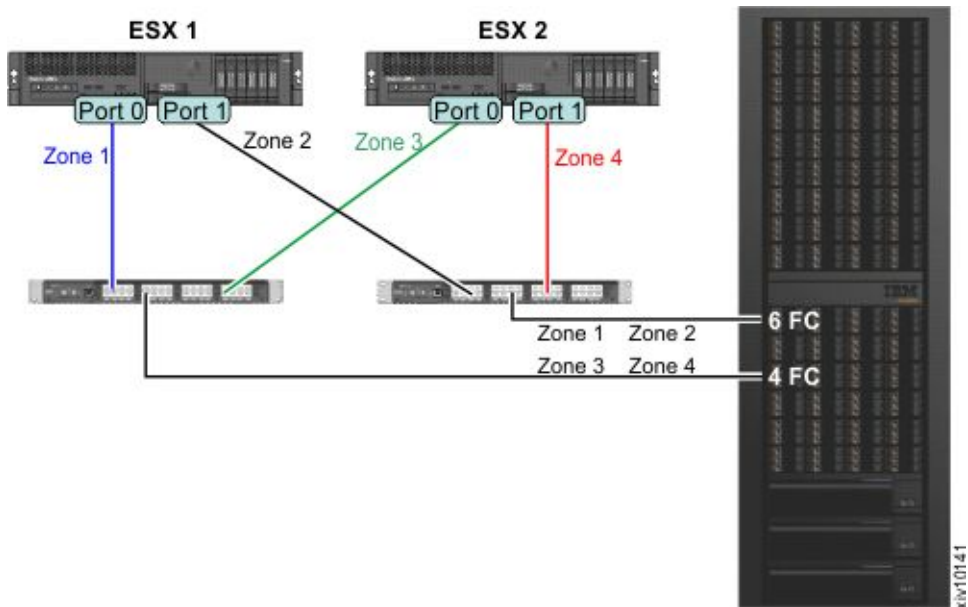


Figure 56. SAN Fibre Channel zoning layout

The following table specifies the devices that are hidden from each other:

Table 16. Initiators and targets in the SAN Fibre Channel zoning layout

| Zone name | Initiator  | Targets                                |
|-----------|------------|----------------------------------------|
| Zone 1    | ESX1_Port0 | XIV_Module4_Port1<br>XIV_Module6_Port1 |
| Zone 2    | ESX1_Port1 | XIV_Module4_Port3<br>XIV_Module6_Port3 |
| Zone 3    | ESX1_Port0 | XIV_Module4_Port1<br>XIV_Module6_Port1 |
| Zone 4    | ESX1_Port1 | XIV_Module4_Port3<br>XIV_Module6_Port3 |

## iSCSI target discovery for ESXi hosts

Part of the iSCSI protocol functionality is discovering which targets are accessible to the initiator.

The VMware iSCSI software initiator offers two methods to discover iSCSI targets:

- **Send Targets** (Dynamic Discovery) – When using the send targets method, the software iSCSI initiator is given one IP address or a list of IP addresses. The initiator then queries each address for available iSCSI targets.
- **Static Discovery** – When using the static discovery method, each iSCSI target's IP address needs to be added manually to the initiator's target list.

When connecting an ESXi server to a storage system using iSCSI, connect at least two NICs to at least two interface modules, resulting in a minimum of four paths for each LUN.

Figure 57 shows two ESXi servers that are connected to a storage system. Each iSCSI initiator has three iSCSI interfaces defined using the **Send Targets** method, resulting in six available paths for each LUN.




Figure 57. Connecting ESXi servers over iSCSI

---

## Multipathing and load balancing with ESXi hosts

VMware ESXi servers provide an I/O multipathing driver that can be used for redundancy and for I/O load balancing between paths.

The IBM storage system is a fully symmetrical Active-Active storage array, in which all LUNs are serviced through all storage ports concurrently. The storage system is presented to ESXi as an Asymmetrical Logical Unit Access (ALUA) array. This does not affect performance or create additional management overhead. ESXi 5.x and ESXi 6.x include native support for the IBM storage system, with **Round Robin** as the default multipathing policy. Because ESXi 5.x and ESXi 6.x identify all paths to the IBM storage system as optimized, no additional configuration is needed.



---

## Notices

These legal notices pertain to IBM Storage Host Software Solutions product documentation.

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

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

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

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd.  
1623-14, Shimotsuruma, Yamato-shi  
Kanagawa 242-8502 Japan*

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

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

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be

incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation  
Attn: Office of Legal Counsel  
650 Harry Road  
San Jose, CA 95120-6099  
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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

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

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.



---

## Trademarks

IBM, the IBM logo, [ibm.com](http://ibm.com), AIX, and the AIX logo are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Copyright and trademark information website ([www.ibm.com/legal/us/en/copytrade.shtml](http://www.ibm.com/legal/us/en/copytrade.shtml)).

HP, HP-UX, and HP UNIX are trademarks or registered trademarks of Hewlett Packard Company in the United States, other countries, or both.

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

Oracle and Solaris are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Veritas is a trademark or registered trademark of Symantec Corporation in the United States and other countries.

VMware, ESX, and ESXi are trademarks or registered trademarks of VMware Corporation in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.



---

# Index

## A

- add ports to a host 46
- admin user privileges
  - user privileges 23
- AIX
  - installing 26
- analyze host configuration 85

## B

- best practices 89

## C

- changes applied on host 97
- choosing ESX configurations 101
- CLI attachment utility 37
- cluster
  - defining from the XIV GUI 55
- cluster definition via XCLI 57
- cluster host attachment 53, 57
  - defining the cluster on XIV 54
  - map private volumes to a specific host in a cluster 61
  - map shared volumes to all hosts in a cluster 60
  - map storage volumes to cluster 59
- cluster installation
  - prerequisites 54
- cluster node attachment 64
- compatibility 2
- concept diagram 1
- configuration changes 65
- connectivity parameters 8
- connectivity types 1

## D

- define a host 43
  - HSM 46, 57
  - XCLI 45
  - XIV GUI 44
- detach the host 33
- devices 70
- diagnostics 83
- disable the HSA 94
- disk on key 33
- display volumes 70
- download site 3

## E

- ESX connectivity 101
- ESX hosts 103
- ESXi hosts 101

## F

- Fibre Channel (FC) connections 6
- Fibre Channel SAN zoning 101
- Fibre Channel zoning
  - zoning 9

## H

- HAK installation package
  - upgrade 31
- host attachment kit 1
- host changes 97
- host configuration
  - display changes 79
- host connectivity
  - monitoring
    - XCLI 89, 90
    - XIV GUI 89, 91
- host definition 43
- host diagnostics 83
- host preparation 22
- host profile 85
- host profiler scheduling 93
- host requirements 22
- host-side acceleration 93, 94, 95
- HP-UX
  - installing 26
- HSA 93, 94, 95

## I

- installation 25
- installation package
  - installing on Windows 28
- introduction 1
- iSCSI
  - daemon configuration 98
  - initiator name 98
  - service 98
- iSCSI connections 7
- iSCSI connectivity
  - CHAP authentication 21
- iSCSI ports
  - configure on XIV storage system 16, 18
- iSCSI settings 15
- iSCSI target discovery 102

## L

- Linux
  - installing 26
- Linux host
  - changes applied 97
- load balancing 99, 103
- local installation
  - performing 26

## M

- map storage volumes 49, 50, 62
- move file system from the HSA 95
- move file system to the HSA 94
- multipath configuration 97
- multipath service(s) 98
- Multipathing 103

## N

- network drive 32
- node attachment 64

## O

- operating system settings 89
- optional host definition 43
- OS requirements 22

## P

- permissions 23
- portable HAK package
  - USB flash drive 33
- portable package
  - using without installing 32
- profile information 85

## R

- requirements 2

## S

- single host attachment 37
- Solaris
  - installing 26
- storage area network
  - boot from SAN 23
- system requirements 22
- systems list 66

## T

- troubleshooting 83

## U

- uninstall
  - software removal 33
  - without detaching the host 35
- uninstall the software 33

## V

- VMware 101

- volume mapping
  - rescan (-R) option 92
- volumes 70

## W

- why 1
- Windows Server host
  - changes applied 99
    - built-in MPIO feature 100
    - disk timeout setting 99
    - iSCSI service 99
    - removal of previous HAK versions 99
    - XIV MPIO Load Balancing agent 99
    - XIV MPIO Load Balancing service 99
  - LUN0 device driver 99
- workflow 3

## X

- XCLI 45
- xiv\_attach 37, 64
- xiv\_devlist 65, 70
  - display additional information 71
- xiv\_diag 83
- xiv\_fc\_admin 65, 79
- xiv\_host\_profiler 85
- xiv\_iscsi\_admin 65, 79
- xiv\_syslist 65, 66





Printed in USA

GA32-1060-13

