

Hitachi Virtual Storage Platform F1500

Hitachi Virtual Storage Platform G1x00

Hitachi Virtual Storage Platform

Hitachi Universal Platform V/VM

SVOS RF 8.2

Mainframe Host Attachment Guide

© 2010, 2018 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara Corporation (collectively "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara Corporation at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara Corporation.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, IBM, Lotus, MVS, OS/390, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Contents

Preface	6
Intended audience.....	6
Product version.....	6
Release notes.....	7
Changes made in this revision.....	7
Referenced documents.....	7
Document conventions.....	8
Conventions for storage capacity values.....	10
Accessing product documentation.....	11
Getting help.....	11
Comments.....	11
Chapter 1: Overview of mainframe operations	12
Mainframe compatibility and functionality.....	12
Connectivity.....	13
CU emulation support.....	16
Program products for mainframe.....	16
System option modes.....	19
Mainframe system function enablement	23
Logical channel subsystems and multiple subchannel sets.....	25
Chapter 2: FICON®/zHPF and ESCON® host attachment	26
Host connection interface specifications.....	26
FICON®/zHPF and ESCON® for zSeries hosts.....	28
Setting Up ESCON® and FICON®/zHPF links	28
Multi-pathing for ESCON® and FICON®/zHPF.....	28
Attaching ESCON® front-end directors.....	29
ESCON® connectivity on zSeries hosts	30
Host adapters, cables, distances, and specifications for ESCON® adapters on zSeries hosts.....	30
Cable lengths and types for ESCON® adapters on zSeries hosts	30
Logical paths and path groups for ESCON® adapters on zSeries hosts.....	31
Data transfer for ESCON® adapters on zSeries hosts	31
Directors and channel extenders for ESCON® adapters on zSeries hosts	31

Identifying the ESCON® storage adapter port ID (SAID) on zSeries hosts..	31
Attaching FICON®/zHPF front-end directors.....	31
FICON®/zHPF features.....	32
FICON®/zHPF physical specifications.....	33
FICON®/zHPF logical specifications.....	34
FICON®/zHPF operating environment.....	35
FICON® and FICON®/zHPF protocol sequence.....	36
Hardware specifications.....	37
FICON® front-end directors.....	37
Logical paths.....	37
Supported topologies.....	39
Point-to-point topology.....	39
Switched point-to-point topology.....	40
Cascaded FICON® topology.....	40
Required high-integrity features.....	42
Viewing path information.....	43
Physical host connection specifications.....	44
Logical host connection specifications.....	46
Enabling zHPF operations.....	47
Activating zHPF feature.....	47
Enabling the zHPF feature.....	48
Installing zHPF in a cascading switch configuration.....	49
Removing the FICON®/zHPF feature.....	49
Chapter 3: Mainframe host operations.....	50
Mainframe configuration.....	50
Storage subsystem IDs (SSIDs).....	50
Mainframe hardware definition.....	51
Hardware definition using IOCP (z/OS®, z/VM®, or z/VSE®).....	51
Hardware definition using HCD (z/OS®).....	54
3990 control unit emulation.....	54
2105/2107 CU emulation: defining CU image, base and alias address range.....	55
Dynamic channel management.....	64
Defining the storage system to z/VM® and z/VSE® systems.....	64
Defining the storage system to z/TPF.....	65
Defining the storage system to mainframe Linux.....	65
Mainframe operations.....	66
Initializing the LVIs.....	66
Device operations: ICKDSF.....	66
z/OS® (MVS) cache operations.....	68
z/VM® (VM/ESA) cache operations.....	71

z/VSE® (VSE/ESA) cache operations.....	71
Chapter 4: zLinux operations.....	73
Overview of zLinux operations.....	73
Running Linux on S/390® or zSeries hosts	74
Attaching FC adapters on zSeries hosts running Linux	74
FCP support.....	74
Accessing storage units using FC protocol from a zSeries host	75
Setting up storage units for zSeries hosts running Linux.....	75
Setting up a Linux system to use FC protocol devices on zSeries hosts.....	76
Adding permanent devices for zSeries hosts running Linux	77
Chapter 5: FICON® and ESCON® migration.....	80
Migrating from ESCON® to native FICON®	80
Native ESCON® configuration.....	80
ESCON® to FICON® migration example for a zSeries or S/390® host	81
FICON® configuration example for a zSeries or S/390® host	82
Migrating from a FICON® bridge to a native FICON® attachment.....	83
FICON® bridge overview.....	83
FICON® bridge configuration	84
FICON® and FICON® channel mixed configuration.....	84
Native FICON® configuration on a zSeries or S/390® host	85
Registered state-change notifications (RSCNs) on zSeries hosts	85
Appendix A: Notes on using zDAC with VSP G1x00 and VSP F1500.....	86
Notes on executing the zDAC function.....	86
Message output.....	86
Specifications of open volumes.....	87

Preface

This document describes and provides instructions for mainframe operations on the Hitachi RAID storage systems. The Hitachi RAID storage system models include the Hitachi Virtual Storage Platform G1000 and G1500 (VSP G1x00), Hitachi Virtual Storage Platform F1500 (VSP F1500), Hitachi Virtual Storage Platform (VSP), and Hitachi Universal Storage Platform V and Hitachi Universal Storage Platform VM (USP V/VM).

Please read this document carefully to understand how to use these products, and maintain a copy for your reference.

Intended audience

This document is intended for system administrators, Hitachi Vantara representatives, and authorized service providers who install, configure, and operate the Hitachi RAID storage systems.

Readers of this document should be familiar with the following:

- Data processing and RAID storage systems and their basic functions.
- The Hitachi RAID storage system and the *Hardware Guide* for the storage system.
- The Hitachi Device Manager - Storage Navigator software and the *System Administrator Guide* for the storage system.
- The mainframe operating systems supported by the Hitachi RAID storage systems.
- The hardware used to attach the Hitachi RAID storage system to the mainframe host, including adapters, switches, and directors.

Product version

This document revision applies to the following versions:

- VSP G1x00, VSP F1500: 80-06-4x or later
- SVOS RF 8.2 or later
- VSP: microcode 70-06-5x or later
- USP V/VM: microcode 60-08- xx or later

Release notes

Read the release notes before installing and using this product. They may contain requirements or restrictions that are not fully described in this document or updates or corrections to this document. Release notes are available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>.

Changes made in this revision

- Added information about the **Soft Fence, SPID Fence, Query Host Access** and **INIT/REFORMAT VERIFYOFFLINE** mainframe system functions ([Mainframe system function enablement \(on page 23\)](#)).

Referenced documents

Virtual Storage Platform G1x00 and F1500 documentation:

- *Provisioning Guide for Mainframe Systems*, MK-92RD8013
- *System Administrator Guide*, MK-92RD8016
- *Hitachi Device Manager - Storage Navigator Messages*, MK-92RD8017
- *Hitachi TrueCopy for Mainframe User Guide*, MK-92RD8018
- *Compatible PAV User Guide*, MK 92RD8026

Virtual Storage Platform documentation:

- *Compatible PAV User Guide*, MK 90RD7012
- *Storage Navigator User Guide*, MK 90RD7027
- *Storage Navigator Messages*, MK 90RD7028
- *User and Reference Guide*, MK 90RD7042

Universal Storage Platform V/VM documentation:

- *Hitachi Compatible PAV for IBM® z/OS® User's Guide*, MK 96RD608
- *Storage Navigator User's Guide*, MK 96RD621
- *User and Reference Guide*, MK-96RD635

Copy Manager for TPF documentation:

- *Administrator's Guide*, MK-92RD129
- *Operations Guide*, MK-92RD131
- *Messages and Codes*, MK-92RD130

IBM® documentation:

- *z Systems Processor Resource/Systems Manager Planning Guide*, SB10 7162-01
- *Storage Subsystem Library, 3990 Transaction Processing Facility Support RPQs*, GA32-0134
- *Implementing Linux with IBM Disk Storage*, SG24-6261
- *IBM System z Connectivity Handbook*, SG24-5444
- *z Systems Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7163

For z/OS® V1 R13:

- *z/OS V1.13 DFSMS Technical Update*, SG24-7961
- *z/OS V1R13.0 DFSMS Advanced Copy Services*, SC35-0428
- *z/OS V1 R13 DFSMSdfp Storage Administration*, SC26-7402-15
- *z/OS V1R13.0 HCD Planning*, GA22-7525
- *z/OS V1R13.0 HCD Users Guide*, SC33-7988

For z/OS® V2 R1:

- *z/OS V2R1.0 DFSMS Advanced Copy Services*, SC23-6847-01
- *z/OS V2R1 DFSMSdfp Storage Administration*, SC23-6860-01
- *z/OS V2R1.0 HCD Users Guide*, SC34-2669-02
- *z/OS V2R1.0 Hardware Configuration Definition (HCD) Planning*, GA32-0907-00

For z/OS® V2 R2:

- *z/OS V2R2.0 DFSMS Advanced Copy Services*, SC23-6847-02
- *z/OS V2R2.0 DFSMSdfp Storage Administration*, SC23-6860-02
- *z/OS V2R2.0 HCD Users Guide*, SC34-2669-03
- *z/OS V2R2.0 Hardware Configuration Definition (HCD) Planning*, GA32-0907-01




Document conventions


This document uses the following typographic conventions:

Convention	Description
Bold	<ul style="list-style-type: none"> ▪ Indicates text in a window, including window titles, menus, menu options, buttons, fields, and labels. Example: Click OK. ▪ Indicates emphasized words in list items.

Convention	Description
<i>Italic</i>	<ul style="list-style-type: none"> Indicates a document title or emphasized words in text. Indicates a variable, which is a placeholder for actual text provided by the user or for output by the system. Example: <pre>pairdisplay -g group</pre> <p>(For exceptions to this convention for variables, see the entry for angle brackets.)</p>
Monospace	Indicates text that is displayed on screen or entered by the user. Example: <code>pairdisplay -g oradb</code>
< > angle brackets	Indicates variables in the following scenarios: <ul style="list-style-type: none"> Variables are not clearly separated from the surrounding text or from other variables. Example: <pre>Status-<report-name><file-version>.csv</pre> Variables in headings.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.

This document uses the following icons to draw attention to information:

Icon	Label	Description
	Note	Calls attention to important or additional information.
	Tip	Provides helpful information, guidelines, or suggestions for performing tasks more effectively.
	Caution	Warns the user of adverse conditions and/or consequences (for example, disruptive operations, data loss, or a system crash).

Icon	Label	Description
	WARNING	Warns the user of a hazardous situation which, if not avoided, could result in death or serious injury.

Conventions for storage capacity values

Physical storage capacity values (for example, disk drive capacity) are calculated based on the following values:

Physical capacity unit	Value
1 kilobyte (KB)	1,000 (10^3) bytes
1 megabyte (MB)	1,000 KB or $1,000^2$ bytes
1 gigabyte (GB)	1,000 MB or $1,000^3$ bytes
1 terabyte (TB)	1,000 GB or $1,000^4$ bytes
1 petabyte (PB)	1,000 TB or $1,000^5$ bytes
1 exabyte (EB)	1,000 PB or $1,000^6$ bytes

Logical capacity values (for example, logical device capacity, cache memory capacity) are calculated based on the following values:

Logical capacity unit	Value
1 block	512 bytes
1 cylinder	Mainframe: 870 KB Open-systems: <ul style="list-style-type: none"> ▪ OPEN-V: 960 KB ▪ Others: 720 KB
1 KB	1,024 (2^{10}) bytes
1 MB	1,024 KB or $1,024^2$ bytes
1 GB	1,024 MB or $1,024^3$ bytes
1 TB	1,024 GB or $1,024^4$ bytes
1 PB	1,024 TB or $1,024^5$ bytes

Logical capacity unit	Value
1 EB	1,024 PB or 1,024 ⁶ bytes

Accessing product documentation

Product user documentation is available on Hitachi Vantara Support Connect: <https://knowledge.hitachivantara.com/Documents>. Check this site for the most current documentation, including important updates that may have been made after the release of the product.

Getting help

[Hitachi Vantara Support Connect](#) is the destination for technical support of products and solutions sold by Hitachi Vantara. To contact technical support, log on to Hitachi Vantara Support Connect for contact information: https://support.hitachivantara.com/en_us/contact-us.html.

[Hitachi Vantara Community](#) is a global online community for Hitachi Vantara customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. **Join the conversation today!** Go to community.hitachivantara.com, register, and complete your profile.

Comments

Please send us your comments on this document to doc.comments@hitachivantara.com. Include the document title and number, including the revision level (for example, -07), and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara Corporation.

Thank you!

Chapter 1: Overview of mainframe operations

This chapter provides an overview of mainframe host attachment, functions, and operations for the Hitachi RAID storage systems.

Mainframe compatibility and functionality

The storage systems provide full System-Managed Storage (SMS) compatibility and support the following functionalities in the mainframe environment:

- Sequential data striping
- Cache fast write (CFW) and DASD fast write (DFW)
- Enhanced dynamic cache management
- Extended count key data (ECKD) commands
- Multiple Allegiance
- Concurrent Copy (CC)
- Peer-to-Peer Remote Copy (PPRC)
- Extended Remote Copy (XRC)
- FlashCopy[®]
- Parallel Access Volume (PAV)
- Hyper Parallel Access Volume (HPAV)
- Priority I/O queuing
- Red Hat Linux for IBM[®] S/390[®] and zSeries[®]
- SUSE Linux for IBM[®] S/390[®] and zSeries[®]
- zHyperWrite for DB2[®] (for details, see *Hitachi TrueCopy[®] for Mainframe User Guide*)
- zHPF Extended Distance II (VSP G1x00 models, VSP F1500 80-05-0x or later)
- FICON[®] Dynamic Routing
- FICON[®] Forward Error Correction (VSP G1x00 models, VSP F1500 80-05-0x or later)
- Soft Fence, SPID Fence (VSP G1x00 models, VSP F1500 80-05-4x or later) (for details, see the *Provisioning Guide for Mainframe Systems*)
- Query Host Access (VSP G1x00 models, VSP F1500 80-05-4x or later) (for details, see the *Provisioning Guide for Mainframe Systems*)

The following table lists the mainframe operating systems currently supported by the storage systems. Please contact your Hitachi Vantara account team for the latest information about mainframe operating system support.

Vendor	Operating systems
IBM®	OS/390® (USP V/VM only)
	MVS/ESA, MVS/XA
	VM/ESA, VSE/ESA (USP V/VM only)
	z/OS® V1, z/OS® V2, z/OS.e, z/VM®, z/VSE®
	TPF, z/TPF
Red Hat	Linux for IBM® zSeries®
Novell	SUSE Linux for IBM® zSeries®
Fujitsu	MSP (USP V/VM only)

Connectivity

The storage systems support all-mainframe, all-open-system, and multiplatform configurations. The front-end director (FEDs) process the channel commands from the hosts and manage host access to cache. In the mainframe environment, the front-end directors perform CKD-to-FBA and FBA-to-CKD conversion for the data in cache.

The front-end directors support the following types of host channel interface: FICON[®]/zHPF, Extended Serial Adapter (ExSA) (compatible with ESCON[®] protocol) (USP V/VM only), and fibre channel (FC). The FICON[®]/zHPF and fibre-channel FED features are available in shortwave (multimode) and longwave (single-mode) versions.

- FICON[®]/zHFP:

The same FICON[®] features (boards) are used for FICON[®] and FICON[®]/zHFP. The FICON[®] features provide data transfer speeds of up to 1,600 MB/sec (16 Gbps) and have 8 or 16 ports per feature (pair of boards). The CHPID must be set as the FC type.

FICON[®] data transmission rates vary according to configuration:

- z800 and z900 series hosts have the following possible configurations:
 - FICON[®] channel will operate at 1 Gbps ONLY.
 - FICON[®] Express channel transmission rates vary according to microcode release. If microcode is 3G or later, the channel will auto-negotiate to set a 1 Gbps or 2-Gbps transmission rate. If microcode is previous to 3G, the channel will operate at 1 Gbps ONLY.
- z9, z10, z12, z114, z196 series hosts support the following configurations:
 - FICON[®] channel will operate at 1/2/4/8 Gbps (by auto-negotiation process).
 - FICON[®] Express4 is supported in z9/z10 series. FICON[®] Express8 is supported in z10 series.
- z13 series hosts support the following configurations:
 - FICON[®] channel will operate at 4/8/16 Gbps (by auto-negotiation process depending on the installed FED type).
 - FICON[®] Express16S is supported by z13/z13s series.
 - FICON[®] Dynamic Routing is supported by z13/z13s series.
- S/390 Parallel Enterprise Servers - Generation 5 (G5) and Generation 6 (G6) only support FICON[®] at 1 Gbps.

IBM[®] FICON[®] channel auto-negotiation rates are:

- FICON[®] Express2 will auto-negotiate to 2 Gbps or 1 Gbps.
- FICON[®] Express4 will auto-negotiate to 4 Gbps, 2 Gbps, or 1 Gbps.
- FICON[®] Express8 will auto-negotiate to 8 Gbps, 4 Gbps, or 2 Gbps.
- FICON[®] Express8S will auto-negotiate to 8 Gbps, 4 Gbps, or 2 Gbps.
- FICON[®] Express16S will auto-negotiate to 16 Gbps, 8 Gbps, or 4 Gbps.

When configured with shortwave (multimode) FICON[®] front-end directors, the storage system can be located according to the following:

- Multimode 62.5 μm : 300m (1 Gbps), 150m (2 Gbps), 70m (4 Gbps), 21m (8 Gbps), 15m (16 Gbps)
- Multimode 50 μm (50/125 μm laser optimized multi-mode (OM4) fiber cable): 860m (1 Gbps), 500m (2 Gbps), 400m (4 Gbps), 190m (8 Gbps), 125m (16 Gbps)

When configured with longwave (single-mode) FICON[®] front-end directors, the storage system can be located according to the following:

- Single-mode 9 μm : 10km, 20km with RPQ

- Fiber-channel: Fiber-channel connectivity is supported for IBM® mainframe attachment when host FICON®/zHPF channel paths are defined to operate in fiber-channel protocol (FCP) mode. The fibre-channel features provide data transfer speeds of up to 8 Gbps and can have 8 or 16 ports per feature (pair of boards). The storage systems support shortwave (multimode) and longwave (single-mode) versions of fibre-channel ports on the same front-end director board.
- ESCON® (USP V/VM only): The ExSA features provide data transfer speeds of up to 17 MB/sec and have 8 ports per feature (pair of boards). Each ExSA channel can be directly connected to a CHPID or a serial channel director. Shared serial channels can be used for dynamic path switching. The USP V/VM also supports the ESCON® Extended Distance Feature (XDF). The VSP G1x00, VSP F1500, and VSP do not support ExSA/ESCON® connectivity.
- FICON® Dynamic Routing (FIDR):
Connections to IBM® z13 or z13s with FICON® 16S channels can utilize FICON® Dynamic Routing for interswitch links on cascaded FICON® director configurations.
The FICON® directors must support the following features:
 - Brocade directors:
 - Exchange Based Routing
 - Control Unit Port (CUP) function
 - Cisco directors:
 - Open Exchange ID Routing
 - Control Unit Port (CUP) function

CU emulation support

The following table specifies the CU emulation support for the storage systems.

Storage system	CU emulation support
Hitachi Universal Storage Platform V/VM	<ul style="list-style-type: none"> ▪ 3990-6, 3990-6E ▪ 2105, 2107
Hitachi Virtual Storage Platform	<ul style="list-style-type: none"> ▪ 3990-6, 3990-6E ▪ 2105, 2107
Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500	<ul style="list-style-type: none"> ▪ 2107

Program products for mainframe

The following tables list and describe the Hitachi software products and the host- and server-based products for use with mainframe systems.

Device Manager - Storage Navigatorr-Based Software for Mainframe Users

Name	Description
Hitachi Dynamic Provisioning	Provides “virtual storage capacity” to simplify storage addition and administration, eliminate application service interruptions, and reduce costs.
Hitachi Dynamic Tiering	Improves Dynamic Provisioning pool storage performance further by taking advantage of performance differences between hardware tiers. Enables you to configure the monitoring frequency modes at which data is accessed and relocate that data in a specific pool and tier based on the results.
Hitachi TrueCopy® for Mainframe	Enables the user to perform remote copy operations between storage systems in different locations. Hitachi TrueCopy® provides synchronous and asynchronous copy modes for mainframe data.
Hitachi ShadowImage® for Mainframe	Allows the user to create internal copies of volumes for purposes such as application testing and offline backup. Can be used in conjunction with TrueCopy to maintain multiple copies of data at primary and secondary sites.
Hitachi Compatible Mirroring for IBM® FlashCopy® V2	Provides compatibility with the IBM® FlashCopy® mainframe host software function, which performs server-based data replication for mainframe data.
Hitachi Universal Replicator software for Mainframe	Provides a storage-based hardware solution for disaster recovery which enables fast and accurate system recovery, particularly for large amounts of data which span multiple volumes. Using UR, you can configure and manage highly reliable data replication systems using journal volumes to reduce chances of suspension of copy operations.
Hitachi Compatible Replication for IBM® XRC	Provides compatibility with the IBM® Extended Remote Copy (XRC) mainframe host software function, which performs server-based asynchronous remote copy operations for mainframe LVIs.
Hitachi Universal Volume Manager	Realizes the virtualization of the storage system. Users can connect other storage systems to the storage system and access the data on the external storage system via virtual devices on the storage system. Functions such as TrueCopy and Cache Residency Manager can be performed on the external data.

Name	Description
Hitachi Virtual Partition Manager	Provides storage logical partition and cache logical partition: <ul style="list-style-type: none"> ▪ Cache logical partition allows you to divide the cache into multiple virtual cache memories to reduce I/O contention. ▪ Storage logical partition (USP V/VM only) allows you to divide the available storage among various users to reduce conflicts over usage.
Hitachi Virtual LVI	Enables users to convert single volumes (LVIs or LUs) into multiple smaller volumes to improve data access performance.
Hitachi Cache Residency Manager	Allows users to “lock” and “unlock” data into cache in real time to optimize access to your most frequently accessed data.
Hitachi Compatible PAV	Enables the mainframe host to issue multiple I/O requests in parallel to single LDEVs. Provides compatibility with the IBM® Workload Manager (WLM) host software function and supports both static and dynamic PAV functionality.
Hitachi Volume Security software	Allows users to restrict host access to data. Mainframe users can restrict host access to LVIs based on node IDs and logical partition (LPAR) numbers.
Hitachi Volume Retention Manager	Allows users to protect data from I/O operations performed by hosts. Users can assign an access attribute to each logical volume to restrict read and/or write operations, preventing unauthorized access to data.

Host/Server-Based Software for Mainframe Users

Name	Description
Hitachi Business Continuity Manager	Enables mainframe users to make Point-in-Time (PiT) copies of production data, without quiescing the application or causing any disruption to end-user operations, for such uses as application testing, business intelligence, and disaster recovery for business continuance.
Hitachi Cross-OS File Exchange Hitachi Code Converter	Enables users to transfer data between mainframe and open-system platforms using the FICON® and/or ESCON® channels, for high-speed data transfer without requiring network communication links or tape.
Hitachi Copy Manager for z/TPF	Enables z/TPF users to control DASD copy functions on storage systems from z/TPF through an interface that is simple to install and use.

Name	Description
Hitachi Cache Manager (USP V/VM only)	Enables users to perform Cache Residency Manager operations from the mainframe host system. Cache Residency Manager allows you to place specific data in cache memory to enable virtually immediate access to this data.
Hitachi Dataset Replication for Mainframe (USP V/VM only)	Operates together with the ShadowImage feature. Rewrites the OS management information (VTOC, VVDS, and VTOCIX) and dataset name and creates a user catalog for a ShadowImage target volume after a split operation. Provides the prepare, volume divide, volume unify, and volume backup functions to enable use of a ShadowImage target volume.

System option modes

To provide greater flexibility, the storage system has additional operational parameters called system option modes (SOMs) that allow you to tailor the storage system to your unique operating requirements. The SOMs are set on the service processor (SVP) by your Hitachi Vantara representative.

- Review the SOMs for your operational environment. The SOMs are described in detail in the System Administrator Guide or Hardware Guide for your storage system.
- Work with your Hitachi Vantara team to make sure the appropriate SOMs are set on your storage system.
- Check each new revision of the System Administrator Guide or Hardware Guide to see if there are any SOM changes that may apply to your operational environment. If so, contact your Hitachi Vantara team.

The following table lists and describes the SOMs that apply to mainframe operations. The SOM information may have changed since this document was published. Contact your Hitachi Vantara team for the latest SOM information.

- SOM: SOM number
- Category: Functions to which the SOM applies
- Description:
 - Function that the SOM provides
 - Default setting (ON or OFF)
 - For remote copy functions, storage system on which the SOM should be set: main control unit (MCU) or remote control unit (RCU)



Caution: Changing the following SOMs is disruptive. Set these SOMs before beginning operations:

- 493 (requires power cycle)
- 494 (requires power cycle)

SOMs for General Mainframe Operations

SOM	Category	Description
305	Mainframe	<p>Enables the pre-label function (creation of VTOC including VOLSER).</p> <ul style="list-style-type: none"> ▪ ON: Prelabel is enabled. ▪ OFF (default): Prelabel is not enabled. <p>Notes:</p> <ol style="list-style-type: none"> 1. Set SOM 305 to ON before performing LDEV Format for a mainframe volume if you want to perform OS IPL (volume online) without fully initializing the volume after the LDEV Format. However, full initialization is required in actual operation. 2. Processing time of LDEV format increases by as much as full initialization takes. 3. The following functions and conditions are not supported: <ul style="list-style-type: none"> ▪ Quick format ▪ 3390-A (Dynamic Provisioning attribute) ▪ Volume Shredder 4. Full initialization is required in actual operation. 5. Microcode versions earlier than 80-05-32-00/00 or earlier than 80-05-47-00/00 (within 80-05-4x range) have a bug. Refer to BID=RAID800-XXXX (SR# HDS04414070). Make sure to use the workaround or the fixed code when SOM 305 is ON.
493	Mainframe USP V/VM only	<p>Controls whether a unique SAID is reported to the host when using the CUIR function.</p> <p>Caution: Power cycle is required when this SOM is changed.</p>

SOM	Category	Description
		<p>The CUIR function requires that the SAID reported to the host is unique. The SAID value cannot be changed during online operations. To change the SAID value from normal to unique, set this SOM to ON and then perform a power cycle. Power cycle includes power OFF/ON (volatile/non-volatile), start-up after breaker OFF/ON, or offline microcode exchange. Setting this SOM to ON without performing a power cycle does not enable the function.</p> <ul style="list-style-type: none"> ▪ ON: When this SOM is ON and a power cycle is performed, a unique SAID value for each port is reported to the host. ▪ OFF (default): When this SOM is OFF and a power cycle is performed and the CU emulation type is 3990 or 2105, normal SAID values are reported to the host. When the CU emulation type is 2107, the SAID value of the mainframe PCB port remains unique, even after setting this SOM to OFF and then performing a power cycle.
494	Mainframe USP V/VM only	<p>Controls whether CUIR processing is used when replacing a FICON[®] PCB.</p> <p>Caution: Power cycle is required.</p> <ul style="list-style-type: none"> ▪ ON: CUIR processing is available when replacing a FICON[®] PCB, but only when SOM 493 is ON and a power cycle is performed to enable the SAID unique mode. Power cycle includes PS-OFF/ON (volatile/non-volatile), start-up after breaker OFF/ON, or offline microcode exchange. ▪ OFF (default): CUIR processing is not available.
531	Mainframe and Open	<p>When PIN data is generated, the SIM currently stored in SVP is reported to the host.</p> <ul style="list-style-type: none"> ▪ ON: The SIM for PIN data generation is stored on the SVP and reported to the host. ▪ OFF (default): The SIM for PIN data generation is stored on the SVP and is not reported to the host.

SOM	Category	Description
545	Mainframe USP V/VM only	<p>When writing to a z/VM[®] Minidisk and creating the record 0 field, this SOM controls whether to allow the record 0 format with the WRFTK (x95) command when the CCHH of the Count part transferred from the host differs from the CCHH of the currently accessed track address.</p> <ul style="list-style-type: none"> ▪ MCU/RCU: This SOM applies only to the MCU. ▪ ON: The record 0 format is allowed. ▪ OFF (default): The record 0 format is not allowed. <p>Note: Set this SOM to ON when the CU emulation type is 2107, or when INVALID TRACK FORMAT ERROR occurs on z/VM[®] MINIDISK.</p>
1091	Mainframe VSP G1x00, and VSP F1500 only	<p>Enables or disables the IBM[®] zHyperWrite function. When the mode setting is changed to ON or OFF, SCI is reported to the host and the zHyperWrite function is enabled or disabled.</p> <ul style="list-style-type: none"> ▪ ON: The zHyperWrite function is enabled. (ReadFeatureCode setting for the zHyperWrite function) ▪ OFF (default): The zHyperWrite function is disabled. <p>Notes:</p> <ol style="list-style-type: none"> 1. The mode can be applied when DKCMAIN version 80-04-21-00/00 or higher that supports the zHyperWrite function is used. 2. To enable the zHyperWrite function, set the mode to ON on both MCU and RCU of TCz. 3. To downgrade the microcode to a version that does not support the zHyperWrite function, set the mode to OFF prior to downgrading the microcode. 4. Even when the mode is set to ON for a storage system in 3DC configuration, the zHyperWrite function does not work for volumes in 3DC configuration. 5. Do not set the mode to ON if SOM474 is set to ON. 6. Do not set the mode to ON if Preserve Mirror configuration is created. 7. Do not set the mode to ON if the TCz Function Switch 12 is set to ON. 8. To use the zHyperWrite function, a HyperSwap environment is required.

SOM	Category	Description
		<p>9. Ensure that the CFW data setting for the PPRC/TCz pair is set to 'Secondary Volume Copy' so that CFW data is copied to the S-VOL. If this is not set, zHyperWrite will not function.</p>

Mainframe system function enablement

In addition to the system option modes (SOMs), additional mainframe system functions can be enabled or disabled by using the Device Manager - Storage Navigator **Edit Mainframe System Functions** window.

The following two functions can be enabled:

- Soft Fence, SPID Fence, Query Host Access
- ICKDSF INIT/REFORMAT VARIFY OFFLINE

Soft Fence, SPID Fence, Query Host Access

Soft Fence and SPID Fence are volume protection functions provided by IBM® that are used for disaster recovery operations. The Query Host Access (QHA) function provides information about DASD and storage control units on IBM® z/OS®. For details about the Soft Fence, SPID Fence, and Query Host Access host functions and related host commands, see the IBM® documentation.

- **Soft Fence:** Soft Fence is defined and canceled from the mainframe host. When you set Soft Fence to pair volumes created by using PPRC commands, the Soft Fence setting prevents unintended access to protect volumes that are needed to recover from a disaster. A volume for which Soft Fence is defined rejects host read/write requests as well as pair operations from the following software products:
 - TrueCopy for Mainframe
 - ShadowImage for Mainframe
 - Universal Replicator for Mainframe
 - Compatible FlashCopy®
 - Compatible XRC
 - Volume Migration
 - FICON® Data Migration

You can check the Soft Fence status (enabled or disabled) of a volume by using the **Mainframe Fenced Status** window in Device Manager - Storage Navigator.

- **SPID Fence:** SPID Fence is defined and canceled from the mainframe host. When you set SPID Fence to a volume, the VARY ONLINE host command is rejected so that the volume is not allowed to be online.

You can check the SPID Fence status (enabled or disabled) of a volume by using the **Mainframe Fenced Status** window in Device Manager - Storage Navigator.

- **Query Host Access:** The Query Host Access function enables you to display the host access information for a volume when you execute the DEVSERV command with the QHA parameter from the host.

For instructions on enabling and disabling the Soft Fence, SPID Fence, Query Host Access function on the storage system and forcibly disabling the Soft Fence and SPID Fence settings on selected volumes, see the *Provisioning Guide for Mainframe Systems*.

ICKDSF INIT and REFORMAT VERIFYOFFLINE support

VERIFYOFFLINE is a parameter for the ICKDSF INIT and REFORMAT commands running on IBM® z/OS®. When you execute these commands with this parameter, if the target volume is made online by another system, the INIT or REFORMAT command cannot be executed. To execute these commands, first execute the VARY OFFLINE command from the system that is using the applicable volume, and then execute the INIT or REFORMAT command again. For details about ICKDSF INIT and REFORMAT VERIFYOFFLINE, see the IBM® documentation.

The INIT/REFORMAT VERIFYOFFLINE function is disabled by default. To use this function, you must enable it on the storage system in advance. Note that this function supports only ports set for I-2107. Even when the INIT/REFORMAT VERIFYOFFLINE function is enabled on the storage system, the function is not supported on ports set for I-2107-TPF.

For instructions on enabling and disabling the INIT/REFORMAT VERIFYOFFLINE function on the storage system, see the *Provisioning Guide for Mainframe Systems*.

Logical channel subsystems and multiple subchannel sets

Logical channel subsystems (LCSS) : Storage products can be connected to any FICON[®], FICON[®] (FCP), or ESCON[®] channel supported by any of the logical channel subsystems implemented on IBM[®] processors. Some IBM[®] processor models can support up to four (4) Logical Channel Subsystems (LCSS0 – LCSS3) that provide the connectivity required by the logical partitions defined in the configuration.

Multiple subchannel sets (MSS): Each logical channel subsystem implemented by IBM[®] processor models can also support up to four (4) subchannel sets (MSS0, MSS1, MSS2, and MSS3). MSS0 is supported by all IBM[®] processor models, while the latest zSeries[®] models also provide support for MSS1 and MSS2. IBM model z13 (MT 2964) supports MSS3.

IBM[®] I/O definitions allow the use of subchannel sets 1, 2, and 3 for use by the subchannel information blocks (SCHIBs) that represent special use devices, such as PAV alias devices or PPRC secondary devices.

For the support provided by your IBM[®] processor model, please consult the latest version of the following IBM[®] manual, and review Appendix F, Machine limits and rules: *Input/Output Configuration Program User's Guide for ICP IOCP, SB10-7163*

Chapter 2: FICON[®]/zHPF and ESCON[®] host attachment

This chapter describes and provides general instructions for attaching the storage system to a mainframe host using a FICON[®]/zHPF or ESCON[®] front-end director.

For details about FICON[®]/zHPF connectivity and supported adapters, switches, and directors for the storage systems, see the Mainframe FICON[®] Support Matrix on the Hitachi Vantara interoperability site: https://support.hitachivantara.com/en_us/interoperability.html

Host connection interface specifications

The following table lists the host connection interface specifications for serial (ESCON[®]) and mainframe fibre (FICON) channels.

Host connection interface specifications

Item	Fibre channel and serial (ESCON [®]) channel
Maximum number of CUs	255
CU address range	0 to FE ^{1,3}
Maximum number of SSIDs	1,020
SSID range	3990-6/6E: 0004 to FFFD 2105 and 2107: 0004 to FEFF
Maximum number of LDEVs	65,280
Number of logical volumes	1 to 65,280 ²
Notes:	
1. The number of CUs that can be connected to one ESCON [®] channel (CHPID) is 16. The CU addresses in the interface with a host are 00 to 0F for ESCON [®] channel.	

Item	Fibre channel and serial (ESCON [®]) channel
<p>The number of CUs that can be connected to one FICON[®] channel (CHPID) is 64 or 255:</p> <ul style="list-style-type: none"> ▪ CU addresses for 2105-F20 emulation and FICON: 00 to 3F ▪ CU addresses for 2107 emulation and FICON: 00 to FE <p>2. The maximum number of logical volumes that can be connected to one FICON[®] channel (CHPID) is 32,768.</p> <p>3. For storage systems operating in IBM[®] TPF mode, the control units that can be accessed are in the range of 00-3F.</p>	

Numbers of logical paths that can be connected

The following table lists the number of logical paths that can be connected for serial and mainframe fibre channels.

Item	Fibre channel	Serial channel
Number of channel ports	8 to 112	8 to 112
Maximum number of logical paths per CU	2,048	2,048
Maximum number of logical paths per port	65,536 ^{1,3} / 261,120 ^{2,3}	512 ⁴
Maximum number of logical paths per front-end director	65,536 ¹ / 261120 ²	1,024
Maximum number of logical paths per storage system	131,072 ¹ / 522,240 ²	8,192

Item	Fibre channel	Serial channel
<p>Notes:</p> <ol style="list-style-type: none"> 1. For 2105-F20 CU emulation. 2. For 2107 CU emulation. 3. The maximum number of paths that can be connected to a host per fibre channel port is 1,024: <ul style="list-style-type: none"> ▪ 2105-F20 emulation: 1,024 host paths × 64 CUs = 65,536 logical paths ▪ 2107 emulation: 1,024 host paths × 255 CUs = 261,120 logical paths 4. The maximum number of logical paths per ESCON® port is 512: <p>32 host paths × 16 CUs = 512 logical paths</p> <p>The maximum number of LDEVs that can be connected per ESCON® channel port is 1,024. Therefore, when 256 devices are defined per CU, the number of logical paths is 128:</p> <p>32 host paths × 4 CUs = 128 logical paths</p> 		

FICON®/zHPF and ESCON® for zSeries hosts

This section describes items you should consider before you configure your system with FICON®/zHPF or ESCON® adapters for zSeries hosts.

Setting Up ESCON® and FICON®/zHPF links

If a system requires x ESCON® links, (where x is the number of links) to provide throughput performance and availability, you can consider consolidating a number of these ESCON® links onto FICON®/zHPF links.

For example, you can map four ESCON® links to a single FICON®/zHPF link and maintain approximately equivalent performance. If the ESCON® channel use is low, you can map six or eight ESCON® links to a single FICON®/zHPF link. If the ESCON® channel use is high, you can map two ESCON® links to a single FICON®/zHPF link.

Multi-pathing for ESCON® and FICON®/zHPF

Operational differences between ESCON® and FICON® technology determine the number of channel path connections that can be utilized to access a storage system.

Device unit addressing limitations of ESCON® (1024) vs. FICON® (32,768) determine how many channel paths are required to address and access the control unit images (LSS) and devices defined on the storage system.

Using ESCON® paths, the storage system can use the 'Multi-Path Reconnect' capability to provide increased IO multiplexing.

FICON[®] does not use the 'Multi-Path Reconnect' capability of ESCON[®]. The higher speed of data transfer and full duplex (bi-directional data flow) capabilities require that an IO operation initiated on a FICON[®] channel to complete on the same FICON[®] channel.

Multiple paths (multi-pathing) is limited by zArchitecture channel subsystem implementations to 8 channel paths maximum per logical control unit.

Multiple paths connected to the same control unit images on a storage system should consist of the same channel path types. While a mixture of ESCON[®] and FICON[®] channels can be defined in HCD or IOCP control unit definitions, these mixed channel configurations should only be used for temporary migration situations.

Hitachi Vantara recommends that a minimum of two channel paths be defined to ensure availability. The actual number of paths to use is determined by workload throughput requirements.

Attaching ESCON[®] front-end directors

You can attach the USP V/VM storage system to a host system using ESCON[®] adapters. The USP V/VM can be configured with up to 96 ESCON[®] physical channel interfaces. The VSP G1x00, VSP F1500 and VSP do not support ESCON[®] connectivity.

ESCON[®] channels have a control unit image addressing limitation of X'0' to X'F'. A USP V/VM storage system can assign multiple ESCON[®] channel adapter paths to be able to connect to any of the 16 control unit groups supported. This provides ESCON[®] channel connectivity to all control units supported.

At installation time, ESCON[®] channels can be assigned to a specific range of real control units. To the ESCON[®] channel, the control units will appear as control unit addresses X'0' to X'F' and defined in HCD/IODF and IOCP with the CUADD parameter specifying a hexadecimal character in the range of X'0' to X'F'.

ESCON[®] channel adapters can also be used for TrueCopy/PPRC remote links.

When ESCON[®] channels are configured as remote copy links, the ESCON[®] channels on the primary USP V/VM storage system are configured as normal ESCON[®] channel paths typically referred to as logical channel ports (LCPs). The ESCON[®] channel adapter ports on the remote USP V/VM are configured as remote channel ports (RCPs).

The preferred method for remote copy links is to use fibre-channel connectivity due to greater data throughput capability.

The control unit images (LSS) support the following characteristics for ESCON[®] channel connectivity:

- 1 to 256 devices (bases and aliases) per control unit image. Each unit address can be defined as a logically emulated disk device and can represent a non-PAV device (such as for devices used by Cross-OS File Exchange), a PAV base device, or a PAV alias unit address.
- 1 to 4,096 devices (bases and aliases) with 16 control unit images
- 1 to 512 logical paths per control unit image
- 1 to 64 logical paths per ESCON[®] port (shared by all control unit images)
- 256 path groups per control unit image
- The USP V supports up to 32 host adapters providing a maximum of 64 ESCON[®] links per machine.
- The USP V/VM storage system supports up to 16 host adapters that provide a maximum of 32 ESCON[®] links per machine.

ESCON[®] connectivity on zSeries hosts

This section describes how to attach a storage unit with ESCON[®] adapters.

The USP V/VM storage system can be attached through ESCON[®] channel paths and ESCON[®] directors to different central-electronic complexes and logical partitions (LPARs).

For optimum availability, you should spread the host or ESCON[®] director connections to different ESCON[®] channel adapter ports over both clusters.

Host adapters, cables, distances, and specifications for ESCON[®] adapters on zSeries hosts

Due to ESCON[®] control unit addressing restrictions, each ESCON[®] channel adapter port can only access 16 control unit images.

Each ESCON[®] host adapter provides two physical channel port connections.

Cable lengths and types for ESCON[®] adapters on zSeries hosts

Maximum distance for ESCON[®] cables is 2 km (1.2 mi) with a 50- μ m multimode fibre and 3 km (1.9 mi) with 62.5- μ m multimode fibre. This distance can be extended 103 km (64 mi), although this is not a recommended practice. These cables are used for connecting the USP V/VM storage system to the following components:

- Host channel controller (direct connect, host to storage system)
- ESCON[®] director (switched)
- ESCON[®] channel extender (switched and XRC)
- Peer controller for remote links connection (TrueCopy/PPRC)

The S/390[®] or zSeries host system attaches to one port of an ESCON[®] host adapter in the storage unit. Each storage unit adapter card has two ports.



Note: Appropriate retention hardware to support cable attachments that control bend-radius limits comes with each ESCON[®] host attachment.

Logical paths and path groups for ESCON[®] adapters on zSeries hosts

Each ESCON[®] adapter card supports two ESCON[®] ports. Each port supports 64 logical paths. At the maximum of 32 ESCON[®] ports, the number of logical paths is 2,048. Each controller image supports up to 64 path groups. Each path group can have up to eight logical paths. Each controller image supports a maximum of 512 logical paths per control unit image.

Data transfer for ESCON[®] adapters on zSeries hosts

The ESCON[®] host adapter supports all data input buffer sizes up to 256 bytes. During write operations, the host adapter requests the minimum hexadecimal pacing count of 'X'02'. For commands whose parameter data length is not determined by the parameter data itself, the full transfer count in the command frame is requested in the first data request. The adapter supports an NDF-R count of 7 (that is, a maximum of eight data requests).

Directors and channel extenders for ESCON[®] adapters on zSeries hosts

The USP V/VM storage system supports IBM[®] ESCON[®] directors 9032 Models 1, 2, 3, and 5. The USP V/VM also supports IBM[®] 9036 channel extenders to the distances that are allowed by the 9036 as described in [Cable lengths and types for ESCON adapters on zSeries hosts \(on page 30\)](#). The storage unit supports the 9729 Wavelength Division Multiplexer channel extender up to 50 km (31 mi).

Identifying the ESCON[®] storage adapter port ID (SAID) on zSeries hosts

ESCON[®] channel path information can be displayed by using the IBM[®] processors HMC console. IO problem determination frames on the HMC can be used to view the connection status and the node identifier as provided by the USP V/VM channel adapter port connection. Information provided by the node identifier identifies the USP V/VM storage system serial number, model, plant of manufacture, and 'TAG' information. The TAG information field provided to the IBM[®] ESCON[®] channel identifies the USP V/VM storage adapter port identification (SAID).

Attaching FICON[®]/zHPF front-end directors

FICON®/zHPF uses a mapping layer technique based on fibre-channel and multiplexing technologies to improve data-transfer rates, reduce the number of channel connections, and support greater distances between the storage systems and mainframe hosts. FICON® and FICON®/zHPF have identical hardware specifications, but FICON®/zHPF offers improved data transfer performance.

Supporting full-duplex data transfers and allowing greater throughput rates over longer distances than ESCON®, FICON®/zHPF delivers faster data access, higher transaction rates, and higher data movement speeds between mainframe servers and storage devices. Infrastructure management is simplified and allows geographically dispersed data centers for optimal business continuity. This simplifies data center operations and reduces hardware and storage management costs.

In summary, FICON®/zHPF provides the following key features:

- Full-duplex, bi-directional data transfers
- Multiple I/O processes
- Allows small and large data transfers to be transmitted at the same time
- High-bandwidth transfer (up to 16 Gbps)
- Greater throughput rates over longer distances than ESCON®
- Lowered interlock between the disk controller and channel
- Executing pipeline CCW
- Use of FCP as the fibre-channel link transport for zHPF

The following sections describe the specifications for FICON®/zHPF.

FICON®/zHPF features

The following table lists the FICON®/zHPF features and specifies storage system support.

Feature	Supported		
	VSP G1x00, VSP F1500	VSP	USP V/VM
Extended Distance FICON®	Yes	Yes ¹	No
z/OS® Discovery and Auto Configuration (zDAC)	Yes	Yes ²	No
Single-track operations	Yes	Yes	Yes
Multi-track operations	Yes	Yes ²	No
Bidirectional function	Yes ³	No	No
QSAM/BSAM/BPAM	Yes ³	No	No

Feature	Supported		
	VSP G1x00, VSP F1500	VSP	USP V/VM
Format write operations	Yes ³	No	No
List Prefetch function	Yes ³	Yes ⁴	No
Notes: <ol style="list-style-type: none"> 1. 70-03-32-00/00 or later 2. 70-04-01-00/00 or later 3. 80-02-01-00/01 or later 4. 70-06-30-00/00 or later – Engineering approval is required to use. 			

FICON®/zHPF physical specifications

The following table compares the physical specifications between ESCON® and FICON® / zHPF.

Item	ESCON® (USP V/VM)	FICON®/zHPF
Physical Interface	ESCON®	ANSI FC-PH/PC-PI-2 etc. For details, see ANSI standard "Fibre Channel Single-Byte Command Code Sets Mapping Protocol – 5".
Logical Interface	ESCON®	ANSI FC-SB-2/FC-SB-3/FC-SB-4/FC-SB-5
Bandwidth and Direction	200 Mbps Unidirectional	1 Gbps / 2 Gbps / 4 Gbps / 8 Gbps / 16 Gbps Bidirectional
Cable	Multi Mode/62.5µm	Single Mode/9µm Multi Mode/50 or 62.5µm
Distance		
Long wave	---	Single Mode/9µm 10 km, 20 km (with RPQ)

Item	ESCON® (USP V/VM)	FICON®/zHPF
Short wave	3 km	Multi Mode/62.5µm 300m (1Gbps), 150m (2Gbps), 70m (4Gbps), 21m (8Gbps), 15m (16 Gbps) Multi Mode/50µm 860m (1Gbps), 500m (2Gbps), 400m (4Gbps), 190m (8Gbps), 125m (16 Gbps)
Connector	ESCON® connector	LC-duplex (see the figure below)
Topologies	Point-to-Point Switched Point-to-Point	Point-to-Point Switched Point-to-Point



FICON®/zHPF logical specifications

The following table compares the logical specifications of FICON®/zHPF and ESCON®.

Item	ESCON®	FICON®/zHPF
CCW handling	Handshaking	Pipelining Protocol interlock is reduced.
Number of concurrent I/Os	Each I/O executes once per channel	Multiple I/Os execute concurrently per one channel
Addressing	1,024 devices per channel	32,768 devices per channel
Data transfer	Half-duplex	Full-duplex

Connectivity improvements in FICON®/zHPF:

- The interlock between the host and RAID is lowered in terms of the protocol.
- Throughput is increased using multiple I/O process and bi-directional data transfer.
- Increases in the number of the connectable devices per channel.

FICON®/zHPF operating environment

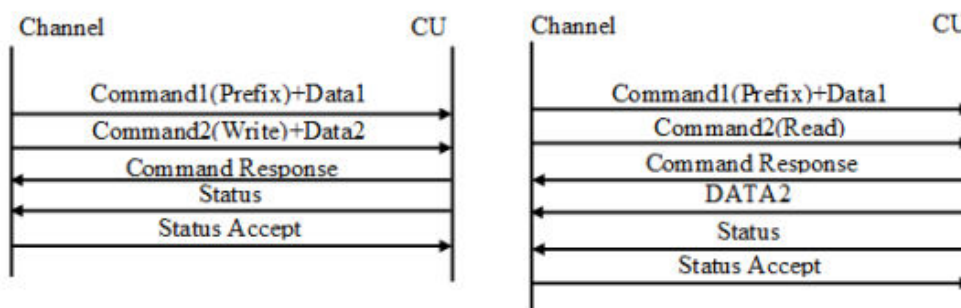
The following table lists the operating environment required for FICON®/zHPF. For more information about FICON®/zHPF operating environments, including supported hosts and operating environments, refer to the IBM® Redbooks.


Item	Specification	
CPU	z9, z10	FICON® Express FICON® Express2 FICON® Express4 FICON® Express8
	z800, z900	Native FICON
		FICON® Express
	z890, z990	FICON® Express
	G5/G6	Native FICON
	z196, z114 zEC12, zBC12	FICON® Express4 FICON® Express8 FICON® Express8Sa
	z13, z13s	FICON® Express16S
Operating system	OS/390 Rev. 2.6 and later releases	
	For FICON®: z/OS® R1.V1 and later releases For FICON®/zHPF: z/OS® V1.7 with the IBM® Lifecycle Extension for z/OS® V1.7 (5637-A01) z/OS® V1.8, V1.9 with PTFs for z10 processor support z/OS® V2.10 and later releases VM/ESA 2.3.0 with appropriate support APARs VM/ESA 2.4.0 and later releases z/VM® Version 3.1.0 and later releases z/VM® Version 4.3.0 and later releases	

Item	Specification
	z/VM® Version 5.4 and later releases z/VM® Version 6 Release 1 and later releases z/VM® Version 6 Release 2 with APAR VM65041 for guest zHPF support
	Linux - (Red Hat / SUSE) With kernel 2.6.38 (released 05-19-2011) provides support for zHPF in LPAR mode and under z/VM® when z/VM® APAR VM65041 is installed.
	TPF V4.1 and later releases z/TPF 1.1 and later releases
	z/VSE® V4.1 and later releases z/VSE® V4.2 and later releases z/VSE® V4.3 and later releases z/VSE® V5.1 and later releases
Control unit emulation	2105-F20 and later 2107
Host-based software	High Performance FICON® (HPF)

FICON® and FICON®/zHPF protocol sequence

The following figures show the basic read/write sequence of FICON® and zHPF Protocols. For details, see the ANSI standard document "Fibre Channel Single-Byte Command Code Sets Mapping Protocol - 5".





Hardware specifications

For details about FICON[®]/zHPF connectivity, FICON[®]/Open intermix configurations, and supported adapters, switches, and directors for the Hitachi RAID storage systems, please contact your Hitachi Vantara account team.

FICON[®] front-end directors

The FICON[®] front-end directors process the channel commands from the hosts and manage host access to cache. In the mainframe environment, front-end directors perform count-key data (CKD)-to-fixed-block architecture (FBA) and FBA-to-CKD conversion for the data in cache. Front-end director boards are installed in pairs.

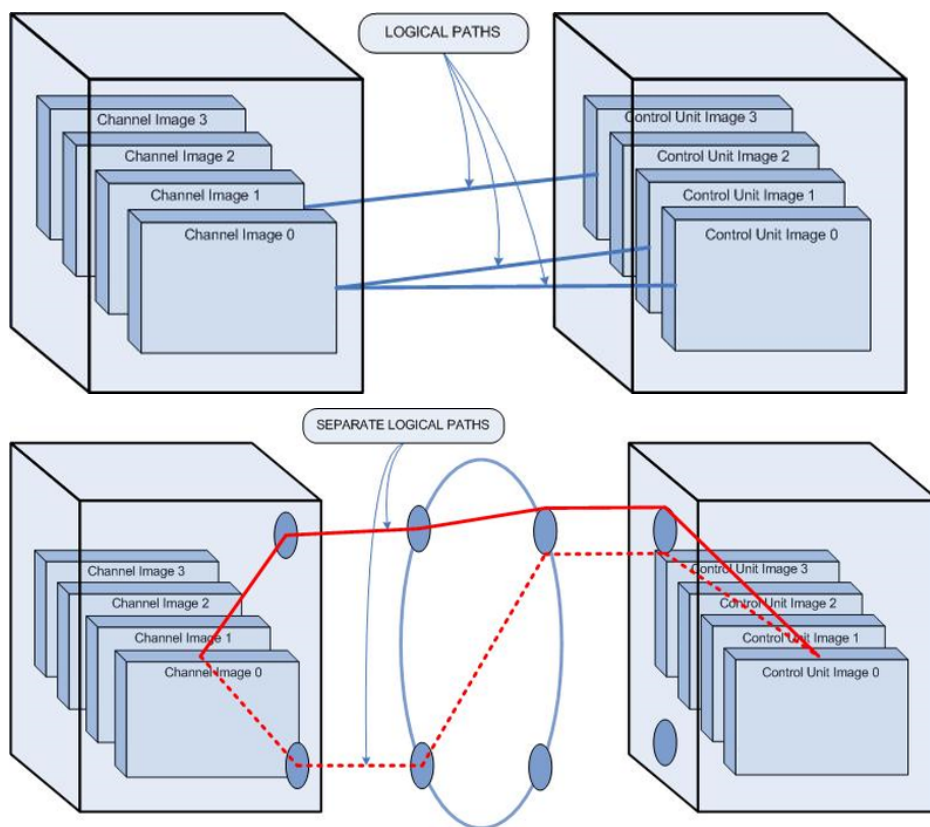
The Hitachi FICON[®] front-end directors provide the storage systems with FICON[®]/zHPF capabilities. FICON[®] front-end directors are available in both shortwave (multimode) and longwave (single-mode) versions. For details about the FICON[®] front-end directors, see the following documents:

- *Hardware Guide for Hitachi Virtual Storage Platform G1000, G1500, and F1500*
- *Hitachi Virtual Storage Platform User and Reference Guide*
- *Hitachi Universal Storage Platform V/VM User and Reference Guide*

Logical paths

FICON[®]/zHPF channels are capable of fast I/O rates. However, mainframes have a limited number of paths that can be defined for a single control unit. One way to address this limitation is to use logical paths. With multiple logical paths, additional paths are defined to the logical control units, providing the host with access to a greater number of control unit interfaces and greater bandwidth. FICON[®] and FICON[®]/zHPF support identical logical paths.


Logical paths are identified through a combination of channel image number, CU image number, channel port address and CU port address. The following figures show two conceptual drawings of mainframe logical paths.



The Hitachi FICON®/zHPF front-end directors provide logical path bindings that map user-defined FICON®/zHPF logical paths. Specifically:

- Each CU port can access 255 CU images (all CU images of logical subsystem).
- To the CU port, each logical host path (LPn) connected to CU port is defined as a channel image number and a channel port address. The LPn identifies the logical host path by the channel image number and channel port address (excluding CU image number).

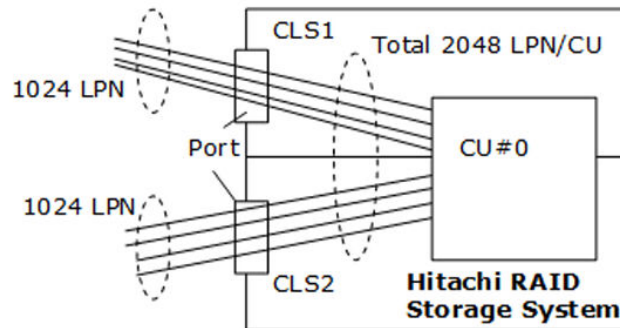
In the following figure, Host Path(1) and Host Path(2) have the same channel image number and CU image number. However, they correspond to different paths because the channel address (and LPn) is different.



The following figure shows another example of logical paths. Instead of being controlled by physical ports, LPNs on the Hitachi RAID storage system are controlled by CU images. Separating LPNs from hardware provides flexibility that allows CU ports to share logical path resources as needed.

Specifically:

- Maximum number of LPN/CU images per storage system = 2,048
- Maximum number of CU images per storage system = 255
- Maximum number of logical paths per storage system = 522,240 (2,048 × 255)




Supported topologies

FICON[®] and FICON[®]/zHPF support the same topologies.

Point-to-point topology

A channel path that consists of a single link interconnecting a FICON[®] channel in FICON[®] native (FC) mode to one or more FICON[®] control unit images (logical control units) forms a point-to-point configuration. A point-to-point configuration is permitted between a channel and control unit when a single control unit is defined on the channel path or when multiple control unit images (logical control units) share the same N Port in the control unit. A FICON[®] channel point-to-point configuration is shown in the following figure.



Switched point-to-point topology

A FICON® channel in FICON® native mode connects one or more processor images to a fibre-channel link, which connects to a fibre-channel switch, and then dynamically to one or more FC switch ports (internally within the switch). From the switch ports, another fibre-channel link to FICON® CU ports interconnects with one or more control units and/or images (logical control units). This forms a switched point-to-point configuration, as shown in the following figure.


Multiple channel images and multiple control unit images can share the resources of the fibre-channel link and fibre-channel switch, permitting multiplexed I/O operations. Channels and control unit links can be attached to the fibre-channel switch in any combination, depending on configuration requirements and available fibre-channel switch resources.

Sharing a control unit through a fibre-channel switch allows communication between a number of channels and the control unit to occur either:

- Over one switch-to-CU link, such as when a control unit has only one link to the fibre-channel switch, or
- Over multiple-link interfaces, such as when a control unit has more than one link to the fibre-channel switch.


In a FICON® switched point-to-point topology, one fibre-channel link is attached to the FICON® channel. From the switch, the FICON® channel communicates with a number of FICON® CUs on different switch ports. At the control unit, the control unit and device-addressing capability is the same as the point-to-point topology. However, communication and addressing capabilities are increased for the channel when it is connected to a fibre-channel switch, and can use the domain and port address portion of the 24-bit N_Port address to access multiple control units.

The communication path between a channel and a control unit consists of the physical channel path and the logical path. In a FICON® switched point-to-point topology with a single switch, the physical paths are the fibre-channel links (interconnection of two fibre-channel links connected by a fibre-channel switch) that provide the physical transmission path between a channel and a control unit.



Cascaded FICON® topology

A FICON® channel in FICON® native (FC) mode connects one or more processor images to a fibre-channel link, which connects to a fibre-channel switch, and then dynamically through one or more FC switch ports (internally within the switch) to a second FC switch in a remote site via FC links. From the second FC switch port, a number of fibre-channel links connect to FICON® CU ports, which interconnect with one or more control units and/or images (logical control units). This forms a cascaded FICON® topology, as shown in the following figure. The cascaded FICON® directors topology is supported only by IBM® zSeries processors.



In a cascaded FICON® topology, one or two switches reside at the topmost (root) level, between the channel (CHL) and disk controller (DKC) ports, as shown in the following figure. With this configuration, multiple channel images and multiple control unit images can share the resources of the fibre-channel link and fibre-channel switches so that multiplexed I/O operations can be performed. Channels and control unit links can be attached to the fibre-channel switches in any combination, depending on configuration requirements and fibre-channel switch resources.

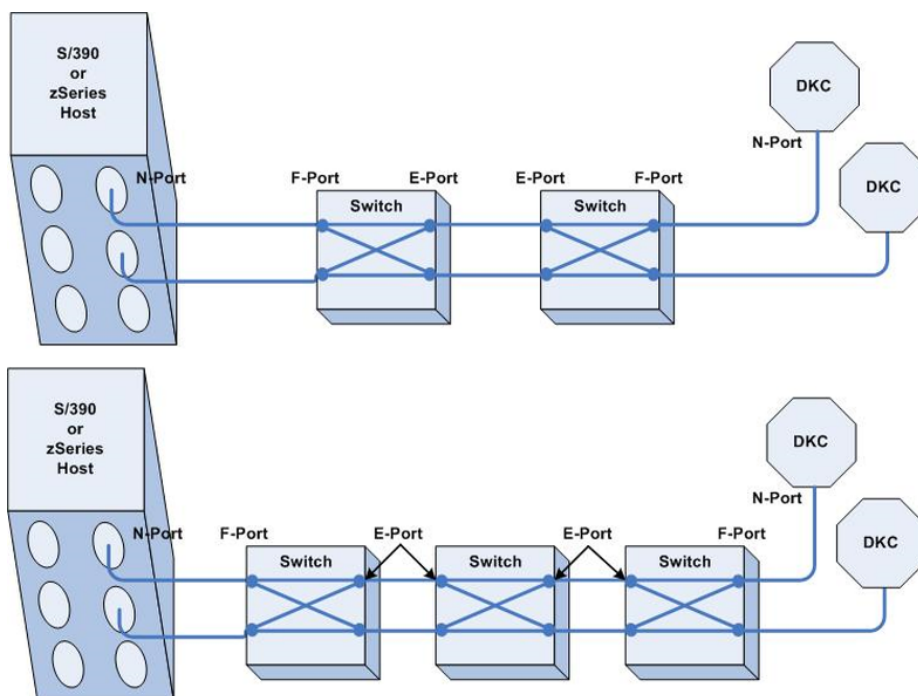
Sharing a control unit through a fibre-channel switch allows communication between a number of channels and the control unit to occur either:

- Over one switch-to-CU link, such as when a control unit has only one link to the fibre-channel switch, or
- Over multiple-link interfaces, such as when a control unit has more than one link to the fibre-channel switch.

In a cascaded FICON® topology, one fibre-channel link is attached to the FICON® channel. From the switch, the FICON® channel communicates with a number of FICON® CUs on different switch ports. At the control unit, the control unit and device-addressing capability is the same as the point-to-point topology. However, communication and addressing capabilities are increased for the channel when it is connected to a fibre-channel switch, and can use the domain and port address portion of the 24-bit N_Port address to access multiple control units.



Note: The domain address portion of the FC 24-bit port address is different, since there are two fibre-channel switches in the channel-to-control unit path.



Note: FICON[®] Dynamic Routing (FIDR) Connections to IBM[®] z13 or z13s with FICON[®] 16S channels can utilize FICON[®] Dynamic Routing for interswitch links on cascaded FICON[®] director configurations. The FICON[®] directors must support the following features:

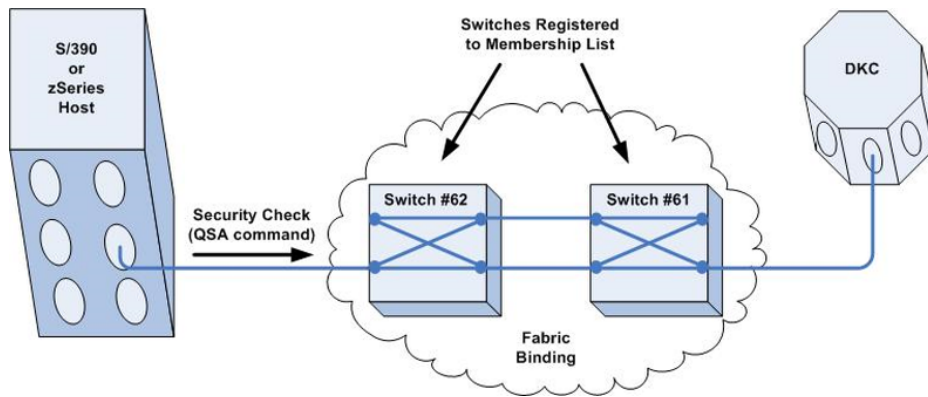
- Brocade directors: Exchange Based Routing, Control Unit Port (CUP) function
- Cisco directors: Open Exchange ID Routing, Control Unit Port (CUP) function

Required high-integrity features

FICON[®] directors in a cascaded configuration require switches that support the following high-integrity features:

- **Fabric binding:** This feature lets an administrator control the switch composition of a fabric by registering WWNs in a membership list and explicitly defining which switches are capable of forming a fabric. In this way, an operator can deny non-authorized switches access into a fabric. A non-authorized switch attempting to gain entry into a fabric becomes isolated and is prevented from accessing fabric resources
- **Insistent Domain IDs:** This feature prohibits the use of dynamic Domain IDs to ensure that predictable Domain IDs are being enforced within the fabric. For example, if a switch has this feature enabled and a new switch is connected to it via an inter-switch links (ISL) without the preferred Domain ID, the new switch is segmented into a separate fabric and user data will not flow. The Insistent Domain IDs feature also ensures that duplicate Domain IDs are not being used within a fabric.

The channel checks whether these high-integrity features are enabled in the switch by issuing a Query Security Attributes (QSA) command during initialization. If these features are not enabled in switch, Channel initialization process stops. See the following figure.



Viewing path information

Depending on the microcode version installed on the storage system, the Link values displayed on the service processor (SVP) Logical Path Status screen may appear as 3-byte values for cascaded topologies:

- The first byte corresponds to the switch address.
- The second byte corresponds to the port address.
- The third byte corresponds to the FC-AL port (this is a constant value in FICON®).

The following figure shows how the link information is displayed in non-cascaded and cascaded topologies.

PT#	NODE	LINK	LGCL
000	HTP1A	26	06
001	HTP1A	26	01
030	HTP1D	27	01
031	HTP1D	27	02

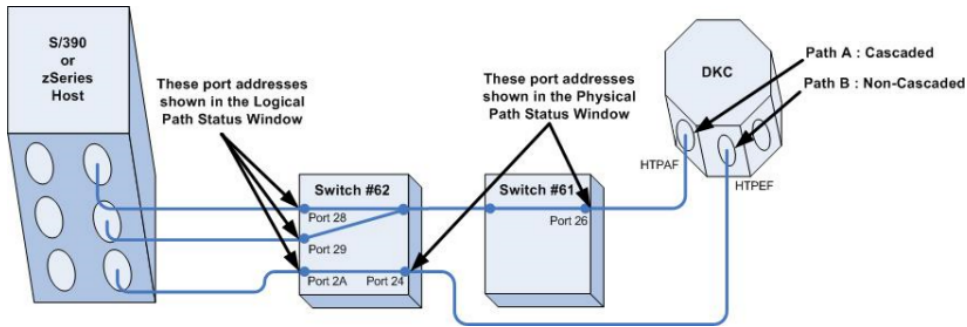
➔

Node	Link(Self)	Link(Dest)	Sta
HTP-1A	640313	FFFFFD	VA
HTP-3A	640413	FFFFFD	VA
HTP-5A	000200	000100	IN
HTP-7A	000000	000000	IN

The following figure shows an example of cascaded and non-cascaded topologies. To confirm the logical path of a FICON® cascade connection in the SVP Maintenance – Mainframe Path – Physical Path Status screen and Logical Path Status screen, note that:

- The first byte in the 3-byte LINK Dest value on the Logical Path Status screen is the address of the switch address at the channel site.
- The first byte in the 3-byte LINK Self value on Physical Path Status screen is the address of the switch at the disk controller (DKC) site.

To confirm the cascade path, compare these two values. The following figure shows an example of a path status display.



The screenshots show the Physical Path and Logical Path status windows. The Physical Path screen displays link details for HTP-AE through HTP-EF. The Logical Path screen displays LGCL values for HTP-FF through HTP-EF. Annotations highlight 'Cascade' and 'Non-Cascade' paths.

Node	Link(Self)	Link(Dest)	Node / Link	LGCL
X HTP-AE	000000	000000	⊗ HTP-FF	1D
X HTP-CE	000000	000000	⊗ HTP-EE	1D
X HTP-EE	612713	661800	⊗ HTP-AF	1E
X HTP-GE	661800	661800	⊗ HTP-EF	1E
X HTP-AF	613613	FFFFFD	⊗ HTP-EF	1D
X HTP-CF	661A13	FFFFFD	⊗ HTP-EF	1E
X HTP-EF	622413	FFFFFD		

path AF
The values of the two switches are different.
➔ Cascade path

path EF
The values of the two switches are the same.
➔ Non Cascade path

Physical host connection specifications

The following table lists the FICON®/zHPF host and RAID physical connection specifications associated with host connections.

	Storage system	Host			
		z800, z900	z890, z990	z9, z10	G5/G6
		Native FICON	FICON® Express	FICON® Express 2 FICON® Express 4 FICON® Express 8	Native FICON
Link bandwidth	1/2/4/8 Gbps (Auto-negotiation) 4/8/16 Gbps (Auto-negotiation)	1 Gbps 4 Gbps	1/2 Gbps 4/8 Gbps	1/2/4/8 Gbps 4/8/16 Gbps	1 Gbps 4 Gbps
Connector	LC-Duplex	SC-Duplex	LC-Duplex	LC-Duplex	SC-Duplex
Cable	Single mode Multi mode	Single mode Multi mode	Single mode Multi mode	Single mode Multi mode	Single mode Multi mode

	Storage system	Host
		z114, z196, zEC12, zBC12, z13, z13s
		FICON® Express4 FICON® Express8 FICON® Express8S FICON® Express16S
Link bandwidth	1/2/4/8 Gbps (Auto-negotiation) 4/8/16 Gbps (Auto-negotiation)	1/2/4/8 Gbps 4/8/16 Gbps

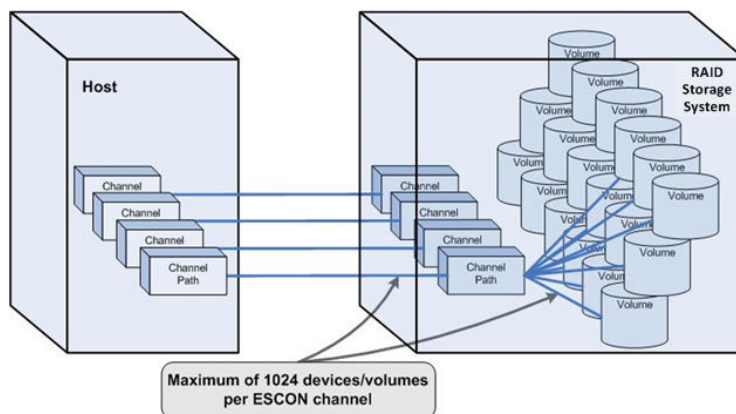
		Host
		z114, z196, zEC12, zBC12, z13, z13s
Storage system		FICON® Express4 FICON® Express8 FICON® Express8S FICON® Express16S
Connector	LC-Duplex	LC-Duplex
Cable	Single mode	Single mode
	Multi mode	Multi mode

Logical host connection specifications


The following figures show examples of two logical host connections. Both examples support ESCON® Multiple Image Facility (EMIF). With ESCON®, the number of Device Addressing per channel port is limited to 1,024. With FICON®/zHPF, the number increases to 32,768.

In the following figure:

- Up to 1,024 unit addresses are supported for each CHL port (limit of the ESCON® CHL port specification). When the number of devices per CHL image is limited to a maximum of 1,024, 16 CU images can be assigned per CHL image. If one CU has 256 devices, the maximum number of CUs per CHL image is limited to 4.
- Up to 32,768 DEVs (all address ranges of the storage system) from one system are supported for each CHL port. This requires that system generation is divided every 1,024 DEVs (4 CUs).



In the following figure, up to 32,768 unit addresses are supported for each CHL port (in this case, the maximum limit is increased). In this example, CHL path configuration can be reduced.



Enabling zHPF operations

Activating zHPF feature

The zHPF license is required to activate the zHPF feature on the storage system. Launch Device Manager - Storage Navigator and use the following information to activate the zHPF feature:

- **License name:** Compatible High Performance Connectivity for IBM® FICON®
- **Permanent with Capacity Free:** The license allows unlimited storage capacity usage on one storage system.

**Caution:**

When FICON® paths are connected to IBM® z13 or IBM® z13s processors using FICON® Express16S channel cards, do not install or remove the High Performance Connectivity for FICON® license.

When the High Performance Connectivity for FICON® license is installed or removed and you are actively connected to an IBM® FICON® Express16S channel type, the logical paths on this channel path will be released.

The logical path is restored by the recovery of the host, but depending on the timing, it is possible that all logical paths of the path group are released, causing access loss to the storage system.

Conditions of occurrence:

- The path is connected to a FICON® Express16S channel on an IBM® z13 or z13s processor.
- The paths are online to the host.
- The High Performance Connectivity for FICON® license is installed or removed.
- The microcode level is:
 - USP V/VM: 60-04-06 or later
 - VSP: 70-02-78-00/00 or later
 - VSP G1x00, VSP F1500: 80-02-24-00/00 or later

Installation or removal of the High Performance Connectivity for FICON® license feature is recommended during a period of low I/O rates or during a maintenance window.

Enabling the zHPF feature

After the zHPF feature has been activated, the zHPF function is not enabled automatically. Use the following procedure to enable the zHPF function.

Procedure

1. You need to enable zHPF manually by performing either Option 1 or Option 2:
 - Option 1: From one device under online CU, issue the **DS QD,Device number,VALIDATE** command. Repeat this operation in all online CUs.
 - Option 2: From one online channel path, issue the **VARY PATH(Device number, Channel path ID),ONLINE** command. Repeat this operation in all online CUs and device numbers.
2. Issue the **D M=DEV(device number)** command to check the zHPF status.. The status **FUNCTIONS ENABLED=ZHPF** indicates that the zHPF function has been enabled.

Installing zHPF in a cascading switch configuration

In the storage system, when the zHPF feature is installed, the zHPF feature is enabled per channel path. For point-to-point (direct) connection and single-switch connection, zHPF is dynamically enabled per channel path. However, in cascading switch connections, it is not automatically enabled: even if you perform option 1 or option 2 above, zHPF IO will not be issued.

Procedure

1. To install the zHPF function in cascading switch configurations, perform either of the following options:
 - Option 1: Vary the channel path in the switch cascading configuration used for zHPF offline with the `CF CHP(Channel path 1-Channel path n),OFFLINE` command, and then vary it online with the `CF CHP(Channel path 1-Channel path n),ONLINE` command. If you always keep the alternate path in the same path group when placing the channel path offline, you can enable zHPF without stopping host access.
 - Option 2: Power cycle the storage system.

Removing the FICON®/zHPF feature

When I/O is running in a z/Linux environment (LPAR, Guest OS), OS may hang if you remove the Compatible High Performance Connectivity for FICON® feature. This is because in the error recovery processing of zHPF I/O that is performed when the storage discontinues the support of zHPF, z/Linux does not have the retry function with the existing I/O (CCW I/O).

Use the following procedure to remove the Compatible High Performance Connectivity for FICON® feature in a z/Linux environment:

Procedure

1. Stop I/Os of all devices under the storage system.
2. Remove the zHPF feature.

Chapter 3: Mainframe host operations

This chapter discusses the operations available for mainframe hosts.

Mainframe configuration

The first step in configuring the storage system is to define the storage system to the mainframe hosts. The three basic areas requiring definition are:

- Storage subsystem IDs (SSIDs)
- Hardware definitions, including I/O Configuration Program (IOCP) or Hardware Configuration Definition (HCD)
- Operating system definitions (HCD or OS commands)



Note: The default primary missing interrupt handler (MIH) time interval that the storage system reports to the host operating system is 30 seconds. Hitachi Vantara recommends that this default MIH time be changed to 45 seconds for those devices that are shared across multiple members of a Sysplex and are not part of a remote copy environment (for example, Universal Replicator, TrueCopy). When devices are part of a remote copy environment, the recommended MIH value is 60 seconds. When devices are the target volumes of a Data Migration effort, the recommended MIH value is 120 seconds.

Storage subsystem IDs (SSIDs)

Storage subsystem IDs (SSIDs) are used for reporting information from the CU (or controller) to the operating system. The SSIDs are assigned by the user and must be unique to all connected host operating environments. Each group of 64 or 256 volumes requires one SSID, so there are one or four SSIDs per CU image. The user-specified SSIDs are assigned during subsystem installation. The following table lists the SSID requirements.

CU emulation	SSID requirements	LVI support	Notes
3990-6, 3990-6E	0004 - FFFD	3390-x, OPEN-x	USP V/VM only.

CU emulation	SSID requirements	LVI support	Notes
2105; 2107*	0004 - FEFF	3390-x, OPEN-x	VSP and USP V/VM support 2105 and 2107. VSP G1x00, VSP F1500 support 2107.
*Compatible PAV operations require that one SSID be set for each set of 256 LDEVs.			

Mainframe hardware definition

Hardware definition using IOCP (z/OS[®], z/VM[®], or z/VSE[®])

The I/O Configuration Program (IOCP) can be used to define the storage system in z/OS[®], z/VM[®], and z/VSE[®] environments (wherever HCD cannot be used). The storage systems support up to 255 control unit (CU) images and up to 65,280 LDEVs. Each CU image can hold up to 256 LDEV addresses. A CU image is the same as an IBM[®] logical sub-system (LSS). The CUADD parameter defines a CU image by its control unit address, which can be 00 to FE for FICON[®]/zHPF connectivity and 0 to F for ESCON[®] connectivity. ESCON[®] connectivity can be used with CU emulation types 3990, 2105, and 2107. FICON[®]/zHPF connectivity can be used with CU emulation types 2105 (2105-F20 or later) and 2107.



Caution: The following are cautions when using IOCP or HCD:

- When multiple LPARs/mainframes can access the volumes, use FEATURE=SHARE for the devices.
- Up to 32,768 addresses per physical interface are allowed by z/OS[®] with FICON[®]/zHPF channels. Multiple groups of FICON[®]/zHPF channels may be used to access the maximum of 65,280 devices that can be configured using 255 control unit images (CUADD=00 through CUADD=FE in the CNTLUNIT statement).
- Up to 1,024 addresses per physical interface are allowed by z/OS[®] with ExSA (ESCON[®]) channels. (This includes PAV base and alias addresses.) ESCON[®] interfaces support up to 4,096 devices for each set of 16 CU images using CUADD=0 through CUADD=F in the CNTLUNIT statement. ESCON[®] channel interfaces can be assigned to access individual groups of 16 CU images.

The following figure shows a sample IOCP definition for FICON® Channels (direct connect and over FICON® switch) of a storage system configured with:

- 2107 CU emulation
- Four FICON®/zHPF channel paths:
 - Two channels paths are connected to a FICON® switch.
 - Two channel paths are directly connected to the storage system.
- Six control unit images (00–05) with 256 LVI per control unit
- 64 base addresses and 128 alias addresses per CU 00, 01, 02, and 03
- 128 base addresses and 128 alias addresses per CU 04 and 05

```

*****
**** Sample FICON CHPID / CNTLUNIT and IODEVICE ****
*****

CHPID PATH=(F9),TYPE=PC,SWITCH=07
CHPID PATH=(FB),TYPE=PC
CHPID PATH=(FD),TYPE=PC
CHPID PATH=(FF),TYPE=PC,SWITCH=07      ^

CNTLUNIT CUNUMR=8100,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=00,UNITADD=(00,256)
CNTLUNIT CUNUMR=8200,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=01,UNITADD=(00,256)
CNTLUNIT CUNUMR=8300,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=02,UNITADD=(00,256)
CNTLUNIT CUNUMR=8400,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=03,UNITADD=(00,256)
CNTLUNIT CUNUMR=8500,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=04,UNITADD=(00,256)
CNTLUNIT CUNUMR=8600,PATH=(F9,FB,FD,FF),      ^
LINK=(04,^^,^^,05),                        ^
UNIT=2107,CUADD=05,UNITADD=(00,256)

IODEVICE ADDRESS=(8100,064),CUNUMR=(8100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8180,128),CUNUMR=(8100),STADET=Y,UNIT=3390A^
,UNITADD=80
IODEVICE ADDRESS=(8200,064),CUNUMR=(8200),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8280,128),CUNUMR=(8200),STADET=Y,UNIT=3390A^
,UNITADD=80
IODEVICE ADDRESS=(8300,064),CUNUMR=(8300),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8380,128),CUNUMR=(8300),STADET=Y,UNIT=3390A^
,UNITADD=80
IODEVICE ADDRESS=(8400,064),CUNUMR=(8400),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8480,128),CUNUMR=(8400),STADET=Y,UNIT=3390A^
,UNITADD=80
IODEVICE ADDRESS=(8500,128),CUNUMR=(8500),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8580,128),CUNUMR=(8500),STADET=Y,UNIT=3390A
IODEVICE ADDRESS=(8600,128),CUNUMR=(8600),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(8680,128),CUNUMR=(8600),STADET=Y,UNIT=3390A

```

The following figure shows a sample IOCP hardware definition for 1,024 LVIs (connected to host CPUs over ESCD) of a USP V/VM storage system configured with:

- 3990 CU emulation
- Two LPARs called PROD and TEST sharing four ExSA (ESCON®) channels connected over two ESCDs to the USP V/VM.

- Each switch has ports C0 and C1 attached to the USP VVM.
- Four (4) CU images (0, 1, 2, 3) with 256 LVIs per control unit.
- Two (2) CU statements per CU image.
- To protect data integrity due to multiple operating systems sharing these volumes, these devices require FEATURE=SHARED.

```

CHPID      PATH=(B0,B1),TYPE=CNC,PARTITION=(PROD,TEST,SHR),SWITCH=01
CHPID      PATH=(B2,B3),TYPE=CNC,PARTITION=(PROD,TEST,SHR),SWITCH=02
CNTLUNIT   CUNUMBR=A000,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=0,UNIT=3990
CNTLUNIT   CUNUMBR=A001,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=0,UNIT=3990
CNTLUNIT   CUNUMBR=A100,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=1,UNIT=3990
CNTLUNIT   CUNUMBR=A101,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=1,UNIT=3990
CNTLUNIT   CUNUMBR=A200,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=2,UNIT=3990
CNTLUNIT   CUNUMBR=A201,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=2,UNIT=3990
CNTLUNIT   CUNUMBR=A300,PATH=(B0,B1),UNITADD=((00,256)),LINK=(C0,C1),CUADD=3,UNIT=3990
CNTLUNIT   CUNUMBR=A301,PATH=(B2,B3),UNITADD=((00,256)),LINK=(C0,C1),CUADD=3,UNIT=3990
IODEVICE   ADDRESS=(A000,256),CUNUMBR=(A000,A001),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A100,256),CUNUMBR=(A100,A101),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A200,256),CUNUMBR=(A200,A201),UNIT=3390,FEATURE=SHARED
IODEVICE   ADDRESS=(A300,256),CUNUMBR=(A300,A301),UNIT=3390,FEATURE=SHARED

```

The following figure shows a sample IOCP hardware definition for 1,024 LVIs (directly connected to CPU) of a USP VVM with:

- 2105 CU emulation
- Eight ExSA (ESCON®) channels directly connected to the USP VVM
- Four control units images with 256 LVIs per control unit
- One CU statement per control unit image
- One hundred twenty-eight (128) 3390 base addresses per CU 0 and 1
- One hundred twenty-eight (128) 3390 alias addresses per CU 0 and 1
- Sixty-four (64) 3390 base addresses in CU 2
- One hundred ninety-two (192) 3390 alias addresses in CU 2
- One hundred twenty-eight (128) 3390 addresses in CU 3
- Sixty-four (64) 3390 base addresses per CU 3
- Sixty-four (64) 3390 alias addresses per CU 3
- To protect data integrity due to multiple operating systems sharing these volumes, these devices require FEATURE=SHARED.



Note:

- If you maintain separate IOCP definitions files and create your SCDS or IOCDS manually by running the IZPIOCP or ICP IOCP program, you must define each CU image on a storage system using one CNTLUNIT statement in IOCP. While it is possible to define a CU image using multiple CNTLUNIT statements in IOCP, the resulting input deck cannot be migrated to HCD due to an IBM® restriction allowing only one CNTLUNIT definition.
- If you receive an error message stating that module IOSTHTC is missing, you can ignore this message. For details, see [Dynamic channel management \(on page 64\)](#).

```

CHPID PATH=(60,61,62,63,64,65,66,67),TYPE=CNC
CNTLUNIT CUNUMBR=8000,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=0,UNIT=2105
CNTLUNIT CUNUMBR=8100,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=1,UNIT=2105
CNTLUNIT CUNUMBR=8200,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=2,UNIT=2105
CNTLUNIT CUNUMBR=8300,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=3,UNIT=2105
IODEVICE ADDRESS=(8000,128),CUNUMBR=(8000),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8080,128),CUNUMBR=(8000),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8100,128),CUNUMBR=(8100),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8180,128),CUNUMBR=(8100),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8200,064),CUNUMBR=(8200),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8240,192),CUNUMBR=(8200),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8300,128),CUNUMBR=(8300),STADET=Y,UNIT=3390,FEATURE=SHARED
IODEVICE ADDRESS=(8380,064),CUNUMBR=(8300),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(83C0,064),CUNUMBR=(8300),STADET=Y,UNIT=3390A,FEATURE=SHARED

```

Hardware definition using HCD (z/OS®)

The Hardware Configuration Definition (HCD) utility can be used to define the storage systems in a z/OS® environment. The HCD procedures for 3990 and 2105 control unit emulation types are described below. FICON®/zHPF interface requires 2105-F20 emulation or 2107.

3990 control unit emulation

To define a USP V/VM with 64 or fewer LVIs, use the same procedure as for an IBM® 3990-6/6E subsystem (see the following table).

Parameter	Value
Control Frame:	
Control unit number	Specify the control unit number.
Control unit type	3990-6 or 3990-6E (using 3990-6 emulation)
Channel path IDs	Specify ExSA or FICON®
Unit address	00 (ExSA or FICON®)
Number of units	64
Array Frame:	
Device number	Specify the first device number.
Number of devices	64
Device type	3390
Connected to CUs	Specify the control unit numbers.

The hardware definition for a USP VVM with more than 64 LVIs is different than that for an IBM® 3990 subsystem (see the following table).

Parameter	Value
Control Frame:	
Control unit number	Specify the control unit number.
Control unit type	NOCHECK ¹ Use UIM 3990 for more than 128 logical paths Use UIM 3990-6 for 128 or fewer logical paths
Channel path IDs	Specify ExSA or FICON®
Unit address	00 (ExSA or FICON®)
Number of units	256
Array Frame:	
Device number	Specify the first device number.
Number of devices	256
Device type ²	3390, 3390B, 3390A, 3390D ³ , 3390S ³
Connected to CUs	Specify the control unit numbers.
<p>Notes:</p> <ol style="list-style-type: none"> 1. The NOCHECK function was introduced by APAR OY62560. Defining the USP VVM as a single control unit allows all channel paths to access all DASD devices. 2. IBM® HCD allows you to use the following device types in HCD definitions: <ul style="list-style-type: none"> ▪ 3390: to define a non-PAV device ▪ 3390B: to define a PAV base device ▪ 3390A: to define a PAV alias device ▪ 3390D: to define a PPRC secondary device ▪ 3390S: to define DB2 backup volumes 3. 3390D and 3390S devices cannot be defined in multiple subchannel set 0 (MSS0). 	

For additional HCD information, please refer to the following manual: *z/OS Hardware Configuration Definition User's Guide*, SC34-2669.

2105/2107 CU emulation: defining CU image, base and alias address range

To define a CU image and the base and alias address range that it will support, use the following example for HCD. The following instructions correspond to the 2105/2107 IOCP definition shown in the following figure.

```

CHPID PATH=(60,61,62,63,64,65,66,67),TYPE=CNC
CNTLUNIT CUNUMBER=8000,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=0,UNIT=2105
CNTLUNIT CUNUMBER=8100,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=1,UNIT=2105
CNTLUNIT CUNUMBER=8200,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=2,UNIT=2105
CNTLUNIT CUNUMBER=8300,PATH=(60,61,62,63,64,65,66,67),      +
          UNITADD=(00,256),CUADD=3,UNIT=2105
IODEVICE ADDRESS=(8000,128),CUNUMBER=(8000),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8080,128),CUNUMBER=(8000),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8100,128),CUNUMBER=(8100),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8180,128),CUNUMBER=(8100),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8200,064),CUNUMBER=(8200),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(8240,192),CUNUMBER=(8200),STADET=Y,UNIT=3390A,FEATURE=SHARED
IODEVICE ADDRESS=(8300,128),CUNUMBER=(8300),STADET=Y,UNIT=3390,FEATURE=SHARED
IODEVICE ADDRESS=(8380,064),CUNUMBER=(8300),STADET=Y,UNIT=3390B,FEATURE=SHARED
IODEVICE ADDRESS=(83C0,064),CUNUMBER=(8300),STADET=Y,UNIT=3390A,FEATURE=SHARED

```



Note:

- The HCD PAV definitions must match the configurations in the storage system. If not, error messages are issued when the hosts are IPL'd or the devices are varied online.
- If you receive an error message stating that module IOSTHTC is missing, you can ignore this message. For details, see [Dynamic channel management \(on page 64\)](#).

Procedure

1. From an ISPF/PDF primary options menu, select the **HCD** option to display the basic **HCD** panel. On this panel you must verify the name of the IODF or IODF.WORK I/O definition file to be used.

```

San Diego OS/390 R2.8 Master MENU
OPTION ==> HC                                SCROLL ==> PAGE

                                USERID - HDS
                                TIME   - 20:23

IS ISMF - Interactive Storage Management Facility
P PDF - ISPF/Program Development Facility
IP IPCS - Interactive Problem Control Facility
R RACF - Resource Access Control Facility
SD SDSF - System Display and Search Facility
HC HCD - Hardware Configuration Definition
BMB EMR BLD - BookManager Build (Create Online Documentation)
EMR EMR READ - BookManager Read (Read Online Documentation)
EMI EMR INDX - BookManager Read (Create Bookshelf Index)
SM SMP/E - SMP/E Dialogs
IC ICSF - Integrated Cryptographic Service Facility
OS SUPPORT - OS/390 ISPF System Support Options
OU USER - OS/390 ISPF User Options
S SORT - DF/SORT Dialogs
X EXIT - Terminate ISPF using list/log defaults

F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIND F6=RCHANGE
F7=UP F8=DOWN F9=SWAP F10=LEFT F11=RIGHT F12=RETRIEVE

```


- On the basic **HCD** panel, select the proper I/O definition file, and then select option 1 to display the **Define, Modify, or View Configuration Data** panel.

```

OS/390 Release 5 HCD
Command ==> _____

                          Hardware Configuration

Select one of the following.

1_ 1. Define, modify, or view configuration data
    2. Activate or process configuration data
    3. Print or compare configuration data
    4. Create or view graphical configuration report
    5. Migrate configuration data
    6. Maintain I/O definition files
    7. Query supported hardware and installed UIMs
    8. Getting started with this dialog
    9. What's new in this release

For options 1 to 5, specify the name of the IODF to be used.

I/O definition file . . . 'HDS.IODF06.WORK'          +
-----
| (C) Copyright IBM Corp. 1990, 1998. All rights reserved. |
-----
F1=Help  F2=Split  F3=Exit  F4=Prompt  F9=Swap  F12=Cancel

```

- On the **Define, Modify, or View Configuration Data** panel, select option 4 to display the **Control Unit List** panel.

```

OS/390 Release 5 HCD
----- Define, Modify, or View Configuration Data -----
Select type of objects to define, modify, or view data.

4_ 1. Operating system configurations
    consoles
    system-defined generics
    EDTs
    esoterics
    user-modified generics
    2. Switches
    ports
    switch configurations
    port matrix
    3. Processors
    partitions
    channel paths
    4. Control units
    5. I/O devices

F1=Help  F2=Split  F3=Exit  F9=Swap  F12=Cancel

```

- On the **Control Unit List** panel, if a 2105/2107 type of control unit already exists, then an 'Add like' operation can be used by inputting an '**A**' next to the 2105/2107 type control unit and pressing **Enter**. Otherwise press **F11** to add a new control unit.

```

Goto Filter Backup Query Help
-----
Control Unit List                               Row 27 of 40
Command ==> _____ Scroll ==> PAGE
Select one or more control units, then press Enter. To add, use F11.

/ CU  Type +      Serial-# + Description
- 3107 SCTC      _____
- 3108 SCTC      _____
- 3109 SCTC      _____
- 310A SCTC      _____
- 4000 3990-6    _____
- 4100 3990-6    _____
- 4200 3990-6    _____
- 4300 3990-6    _____
- 5000 3990      _____
- 5001 3990      _____
- 6000 3990      _____
- 6001 3990      _____
- 7000 3990      _____
- 7001 3990      _____
F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F7=Backward
F8=Forward   F9=Swap       F10=Actions  F11=Add       F12=Cancel

```

5. On the **Add Control Unit** panel, input the following new information, or edit the information if preloaded from an 'Add like' operation, and then press **Enter**:
 - Control unit number
 - Control unit type - 2105 or 2107
 - Switch information only if a switch exists. Otherwise leave switch and ports blank.

```

----- Add Control Unit -----
Specify or revise the following values.
Control unit number . . . . . 8000 +
Control unit type . . . . . 2105_____ +
Serial number . . . . . _____
Description . . . . . new 2105 type for 80xx devices__
Connected to switches . . . _ _ _ _ _ _ _ _ _ _ _ _ +
Ports . . . . . _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ _ +
If connected to a switch, select whether to have CHPIDs/link
addresses, and unit address range proposed.
Auto-assign . . . . . 2 1. Yes
                       2. No
F1=Help  F2=Split  F3=Exit  F4=Prompt  F5=Reset  F9=Swap
F12=Cancel

```

6. On the **Select Processor / Control Unit** panel, input an **S** next to the PROC. ID, and then press **Enter**.

```

Goto Filter Backup Query Help
----- Select Processor / Control Unit -----
Row 1 of 1 More: >
Command ==> _____ Scroll ==> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : 8000      Control unit type . . . : 2105

          Log. Addr. -----Channel Path ID . Link Address + -----
/ Proc. ID Att. (CUADD) + 1---- 2---- 3---- 4---- 5---- 6---- 7---- 8----
s PROD
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F6=Previous  F7=Backward   F8=Forward   F9=Swap        F12=Cancel

F6=Previous  F7=Backward   F8=Forward   F9=Swap        F12=Cancel

```

7. On the **Add Control Unit** panel, enter CHPIDs that attach to the control unit, the control unit address, the device starting address, and the number of devices supported, and then press **Enter**.

```

Goto Filter Backup Query Help
----- Add Control Unit -----

Specify or revise the following values.

Control unit number . . : 8000      Type . . . . . : 2105
Processor ID . . . . . : PROD

Channel path IDs . . . . 60  61  62  63  64  65  66  67  +
Link address . . . . .  _  _  _  _  _  _  _  _  +
Unit address . . . . .  00  _  _  _  _  _  _  _  +
Number of units . . . . 256  _  _  _  _  _  _  _

Logical address . . . . 0_ + (same as CUADD)
Protocol . . . . .  _  + (D,S or S4)
I/O concurrency level . 2 + (1, 2 or 3)

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F9=Swap
F12=Cancel

```

8. Verify that the data is correct on the **Select Processor / Control Unit** panel, and then press **F3**.

```

Goto Filter Backup Query Help
----- Select Processor / Control Unit -----
Row 1 of 1 More: >
Command ==> _____ Scroll ==> PAGE

Select processors to change CU/processor parameters, then press Enter.

Control unit number . . : 8000      Control unit type . . . : 2105

      Log. Addr. -----Channel Path ID . Link Address + -----
/ Proc. ID Att. (CUADD) + 1---- 2---- 3---- 4---- 5---- 6---- 7---- 8----
PROD   Yes  0      60   61   62   63   64   65   66   67
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F6=Previous  F7=Backward   F8=Forward   F9=Swap        F12=Cancel

```

9. On the **Control Unit List** panel, add devices to the new **Control Unit**, input an **S** next to CU 8000, and then press **Enter**.

```

Goto Filter Backup Query Help
----- Control Unit List -----
Row 40 of 41
Command ==> _____ Scroll ==> PAGE

Select one or more control units, then press Enter. To add, use F11.

/ CU  Type +      Serial-# + Description
_ 7001 3990      _____
s 8000 2105      _____ add 2105 type for 80xx devices
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F7=Backward
F8=Forward   F9=Swap       F10=Actions  F11=Add        F12=Cancel

```

10. On the **I/O Device List** panel, press **F11** to add new devices.

```

Goto Filter Backup Query Help
-----
I/O Device List
Command ==> _____ Scroll ==> PAGE
Select one or more devices, then press Enter. To add, use F11.
Control unit number : 8000 Control unit type . : 2105

-----Device----- --#-- -----Control Unit Numbers + -----
/ Number Type + PR OS 1--- 2--- 3--- 4--- 5--- 6--- 7--- 8--- Base
***** Bottom of data *****
_____

F1=Help F2=Split F3=Exit F4=Prompt F5=Reset F7=Backward
F8=Forward F9=Swap F10=Actions F11=Add F12=Cancel

```

11. On the **Add Device** panel, enter the following, and then press **Enter**:

- Device number
- Number of devices
- Device type: 3390, 3390B for PAV base device, or 3390A for PAV alias device.

```

Goto Filter Backup Query Help
-----
Add Device -----
Specify or revise the following values.
Device number . . . . . 8000 (0000 - FFFF)
Number of devices . . . . . 128
Device type . . . . . 3390B +
Serial number . . . . . _____
Description . . . . . 3390 Base addresses 8000-807F__
Volume serial number . . . . . _____ (for DASD)
Connected to CUs . . 8000 _____ +
F1=Help F2=Split F3=Exit F4=Prompt F5=Reset F9=Swap
F12=Cancel

```

12. On the **Device / Processor Definition** panel, add this device to a specific **Processor/System-ID** combination by inputting an **S** next to the **Processor**, and then press **Enter**.

```

----- Device / Processor Definition -----
Row 1 of 1
Command ==> _____ Scroll ==> PAGE
Select processors to change device/processor definitions, then press
Enter.
Device number . . . : 8100          Number of devices . . : 128
Device type . . . : 3390B

/ Processor ID  UA + Time-Out  STADET  CHPID + Preferred Explicit Device
s  PROD         No           Yes     + Candidate List          *****
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F6=Previous  F7=Backward  F8=Forward  F9=Swap       F12=Cancel

```

13. On the **Define Device / Processor** panel, enter the values shown in the figure below, and then press **Enter**.

```

----- Define Device / Processor -----
Specify or revise the following values.
Device number . . : 8000          Number of devices . . . . : 128
Device type . . . : 3390B
Processor ID . . . : PROD

Unit address . . . . . 00 + (Only necessary when different from
                          the last 2 digits of device number)
Time-Out . . . . . No (Yes or No)
STADET . . . . . Yes (Yes or No)

Preferred CHPID . . . . . +
Explicit device candidate list . No (Yes or No)

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset      F9=Swap
F12=Cancel

```

14. On the **Define Processor / Definition** panel, verify that the proper values are displayed, and then press **Enter**.

```

----- Device / Processor Definition -----
Row 1 of 1
Command ==> _____ Scroll ==> PAGE
Select processors to change device/processor definitions, then press
Enter.
Device number . . . : 8000          Number of devices . . : 128
Device type . . . : 3390B

/ Processor ID  UA + Time-Out  STADET  CHPID + Preferred Explicit Device
_  PROD         00 No           Yes     + Candidate List          *****
***** Bottom of data *****

F1=Help      F2=Split      F3=Exit      F4=Prompt      F5=Reset
F6=Previous  F7=Backward  F8=Forward  F9=Swap       F12=Cancel

```

15. On the **Define Device to Operating System Configuration** panel, input an **S** next to the Config ID, and then press **Enter**.

```

----- Define Device to Operating System Configuration -----
| Command ==> _____ Scroll ==> PAGE Row 1 of 1
|
| Select OSs to connect or disconnect devices, then press Enter.
|
| Device number . . : 8100          Number of devices : 128
| Device type . . . : 3390B
|
| / Config. ID  Type   Description          Defined
| s PROD      MVS
| ***** Bottom of data *****
|
| F1=Help      F2=Split   F3=Exit     F4=Prompt   F5=Reset
| F6=Previous  F7=Backward F8=Forward  F9=Swap    F12=Cancel
-----

```

16. The **Define Device Parameters / Features** panel displays the default device parameters.



Note: The WLMPAV parameter defaults to “YES”. Set the desired parameters, and then press **Enter**.

```

----- Define Device Parameters / Features -----
| Command ==> _____ Scroll ==> PAGE Row 1 of 6
|
| Specify or revise the values below.
|
| Configuration ID . . : PROD
| Device number . . . : 8000          Number of devices : 128
| Device type . . . . : 3390B
|
| Parameter/
| Feature  Value  P Req.  Description
| OFFLINE  No     P      Device considered online or offline at IPL
| DYNAMIC  Yes     P      Device supports dynamic configuration
| LOCANY   No     P      UCB can reside in 31 bit storage
| WLMPAV   Yes     P      Device supports work load manager
| SHARED   Yes     P      Device shared with other systems
| SHAREDUP No     P      Shared when system physically partitioned
| ***** Bottom of data *****
|
| F1=Help      F2=Split   F3=Exit     F4=Prompt   F5=Reset
| F7=Backward  F8=Forward F9=Swap    F12=Cancel
-----

```

17. This returns to the **Define Device to Operating System Configuration** panel. Press **F3**.
18. The **Update Serial Number, Description and VOLSER** panel now displays the desired device addresses. To add more control units or device addresses, repeat the previous steps.

```

----- Update Serial Number, Description and VOLSER -----
|                               Row 1 of 128 |
| Command ==> _____ Scroll ==> PAGE |
| Device number . . . : 8000      Number of devices : 128 |
| Device type . . . : 3390B |
| Specify or revise serial number, description and VOLSER. |
| Device Number Serial-# Description VOLSER |
| 8000 _____ 3390 Base addresses 8000-807F _____ |
| 8001 _____ 3390 Base addresses 8000-807F _____ |
| 8002 _____ 3390 Base addresses 8000-807F _____ |
| 8003 _____ 3390 Base addresses 8000-807F _____ |
| 8004 _____ 3390 Base addresses 8000-807F _____ |
| 8005 _____ 3390 Base addresses 8000-807F _____ |
| 8006 _____ 3390 Base addresses 8000-807F _____ |
| 8007 _____ 3390 Base addresses 8000-807F _____ |
| 8008 _____ 3390 Base addresses 8000-807F _____ |
| 8009 _____ 3390 Base addresses 8000-807F _____ |
| 800A _____ 3390 Base addresses 8000-807F _____ |
| F1=Help F2=Split F3=Exit F5=Reset F7=Backward |
| F8=Forward F9=Swap F12=Cancel |
-----

```

Dynamic channel management

The storage systems can be used with the IBM® Dynamic Channel Management (DCM) facility. For storage systems defined to operate in IBM® 2105 control unit emulation, a Control Unit Module Table (CUMT) is required. If the CUMT is not installed and you IPL a storage system on a z/OS® operating system that supports DCM, you will receive the following z/OS® IOS message:

```
IOS357I MANUFACTURER HTC LOAD MODULE IOSTHTC NOT FOUND IN THE LNKLST
```

The CUMT named “IOSTHTC” needs to be assembled and link-edited into a LNKLST library. Please contact your Hitachi Vantara representative to obtain the latest jobstream to create this CUMT.



Note: If you receive an error message during IPL stating that module IOSTHTC is missing, you can ignore this message.

Defining the storage system to z/VM® and z/VSE® systems

64 or Fewer LVIs: To define a storage system with less than or equal to 64 LVIs to a z/VM® system, use the same procedure as for an IBM® 3990-6E subsystem. To define a storage system with more than 64 LVIs to z/VM®, enter the LVI address range, storage type, and sharing option for the subsystem as shown below (the address range varies for each installation).

```
[Address Range] TYPE DASD SHARED YES
```

More than 64 LVIs: To define a storage system with more than 64 LVIs to z/VSE®, use the same procedure as for an IBM® 3990-6E subsystem. For a storage system with more than 64 LVIs, the **ADD cuu:cuu** ECKD statements are the same as for the IBM® 3390.



Caution: There may be APAR requirements at different versions of the z/VM[®] software. In addition, certain requirements apply when z/VM[®] is supporting guests (for example, z/OS[®], OS/390[®]) that can run 2105 native. For further information, please refer to the IBM[®] documentation.

Defining the storage system to z/TPF

The storage systems support the IBM[®] Transaction Processing Facility (z/TPF) and Multi-Path Locking Facility (MPLF) functions in either native mode or under z/VM[®]. MPLF support requires z/TPF or TPF 4.1 Operating System. The zTPF/MPLF capability of the storage system enables high levels of concurrent data access across multiple channel paths.



Caution:

- Hitachi provides specific microcode releases for z/TPF operations. You must contact Hitachi Vantara technical support to obtain the correct microcode releases for your operational environment.
- For storage systems operating in IBM[®] TPF mode, the control units that can be accessed are in the range of 00-3F.

For further information about z/TPF and MPLF operations, please refer to the following IBM[®] documentation:

- *Storage Subsystem Library, 3390 Transaction Processing Facility Support RPQs*, IBM[®] document number GA32-0134
- *Storage Subsystem Library, 3990 Storage Control Reference for Model 6*, IBM[®] document number GA32-0274

Defining the storage system to mainframe Linux

The storage systems support attachment to the following mainframe Linux operating systems:

- Red Hat[®] Linux for S/390[®] and zSeries
- SUSE Linux Enterprise Server for IBM[®] zSeries

For information about supported versions, Linux kernel levels, and other details, contact your Hitachi Vantara representative.

For information about preparing the storage system for Linux host attachment, please refer to the following IBM® publication:

- *Implementing Linux with IBM Disk Storage*, SG24-6261

The following links (current at time of document publication) provide information about the different methods for using Linux on mainframe hosts:

- IBM® Enterprise Storage Server (ESS) library

- Linux on S/390® hosts:

<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg244987.html?Open>

- FCP use on zSeries hosts:

<http://www.redbooks.ibm.com/abstracts/sg246694.html?Open>

Mainframe operations

Initializing the LVIs

The LVIs on the storage systems require only minimal initialization before being brought online. The following example shows a z/OS® ICKDSF JCL example of a minimal init job to write a volume ID (VOLID) and volume table of contents (VTOC).



Note: PAV base and alias devices require additional definition. For details, see the *Hitachi Compatible PAV User Guide* for the storage system.

```
// EXAMPLE JOB
// EXEC PGM=ICKDSF, PARM=NOREPLYU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
INIT UNITADDRESS (XXXX) VFY (*NONE*) VOLID (YYYYYY) -
OWNERID (ZZZZZZ)
/*
```

where XXXX = physical install address, YYYYYY = new volume ID, ZZZZZZ = volume ID owner

Device operations: ICKDSF

The storage systems support the ICKDSF media maintenance utility. The ICKDSF utility can also be used to perform service functions, error detection, and media maintenance. Since the storage system is a RAID device, there are only a few differences in operation from conventional DASD or other RAID devices. The following table lists ICKDSF commands that are specific to the storage systems, as contrasted to RAMAC.

Command	Argument	Storage system	Return code
INSPECT	KEEPIT	RAMAC	CC = 12 Invalid parameter(s) for device type.
		Hitachi RAID	CC = 12, F/M = 04 (EC=66BB).
	PRESERVE	RAMAC	CC = 4 Parameter ignored for device type.
		Hitachi RAID	CC = 12, F/M = 04 (EC=66BB) Unable to establish primary and alternate track association for track CCHH=xxxx.
	SKIP	RAMAC	CC = 4 Parameter ignored for device type - skip.
		Hitachi RAID	CC = 12, F/M = 04 (EC=66BB) Primary track CCHH-xxxx found unrecoverable
	NOPRESERVE, NOSKIP,	RAMAC	CC = 0, ALT information not displayed.
	NOCHECK	Hitachi RAID	CC = 0
	ALLTRACKS, ASSIGN, RECLAIM	RAMAC	CC = 12 Invalid parameter(s) for device type.
	Hitachi RAID	In case of PRESERVE: CC = 12, In case of NO PRESERVE: CC = 0.	
INSTALL	SETMODE (3390)	RAMAC	CC = 0 (but not recommended by IBM).
		Hitachi RAID	CC = 0
	SETMODE (3380)	RAMAC	CC = 12, Invalid parameter(s) for device type.
		Hitachi RAID	CC = 12, Function not supported for nonsynchronous DASD.
ANALYZE		RAMAC	CC = 0
		Hitachi RAID	CC = 0

Command	Argument	Storage system	Return code
BUILDX		RAMAC	CC = 0
		Hitachi RAID	CC = 0
REVAL	REFRESH	RAMAC	CC = 12 Device not supported for the specified function.
		Hitachi RAID	CC = 12, F/M = 04 (EC=66BB) Error, not a data check. Processing terminated.
	DATA, NODATA	RAMAC	CC = 0, Data/Nodata parameter not allowed.
		Hitachi RAID	CC=0
CONTROL		RAMAC	CC = 0, ALT information not displayed.
		Hitachi RAID	CC = 0, ALT information not displayed.
INIT		RAMAC	CC = 0, ALT information not displayed.
		Hitachi RAID	CC = 0
REFORMAT		RAMAC	CC = 0, ALT information not displayed.
		Hitachi RAID	CC=0
CPVOLUME		RAMAC	CC = 0, Readcheck parameter not allowed.
		Hitachi RAID	CC=0

z/OS® (MVS) cache operations

To display the cache statistics for the storage systems under z/OS® DFSMS, use the operator command: **D SMS, CACHE**. The following figure shows the cache statistics reported by the storage systems. The storage systems report cache statistics for each SSID in the storage system. Because the dynamic cache management algorithm has been enhanced, the read and write percentages for the storage systems are displayed as N/A. For further information about z/OS® DFSMS cache reporting, please refer to the IBM® document *DFSMSdfp Storage Administrator Reference* (SC28-4920).

SSID	DEVS	READ	WRITE	HIT RATIO	FW BYPASSES
0004	15	N/A	N/A	50%	0
0005	11	N/A	N/A	0%	0
0006	11	N/A	N/A	87%	0
0007	10	N/A	N/A	87%	0

SSID=SUBSYSTEM IDENTIFIER
 DEVS=NUMBER OF MANAGED DEVICES ATTACHED TO SUBSYSTEM
 READ=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR CACHING
 WRITE=PERCENT OF DATA ON MANAGED DEVICES ELIGIBLE FOR FAST WRITE

HIT RATIO=PERCENT OF READS WITH CACHE HITS
 FW BYPASSES=NUMBER OF FAST WRITE BYPASSES DUE TO NVS OVERLOAD

The storage systems support the following z/OS® cache operations:

- **IDCAMS LISTDATA COUNTS:** When the <system> parameter is used with the LISTDATA command, the user must issue the command once for each SSID to view the entire storage system image. The following is a JCL example of the LISTDATA COUNTS command.

```
//LIST JOB. . . . .
//COUNT1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
LISTDATA     COUNTS     VOLUME (VOL000)     UNIT (3390)     SUBSYSTEM
LISTDATA     COUNTS     VOLUME (VOL064)     UNIT (3390)     SUBSYSTEM
LISTDATA     COUNTS     VOLUME (VOL128)     UNIT (3390)     SUBSYSTEM
LISTDATA     COUNTS     VOLUME (VOL192)     UNIT (3390)     SUBSYSTEM
/*
```


- **Subsystem counter reports:** The cache statistics reflect the logical caching status of the volumes. For the storage systems, Hitachi Vantara recommends that you set the nonvolatile storage (NVS) ON and the DASD fast write (DFW) ON for all logical volumes. This will not affect the way the storage system caches data for the logical volumes.

The default caching status for the storage systems is:

- **CACHE ON** for the storage system
- **CACHE ON** for all logical volumes
- **CACHE FAST WRITE ON** for the storage system
- **NVS OFF** for the storage system: set **NVS ON** for the Hitachi storage systems.
- **DFW OFF** for all volumes: set **DFW ON** for Hitachi storage systems.



Note: In normal cache replacement, bypass cache, or inhibit cache loading mode, the storage systems perform a special function to determine whether the data access pattern from the host is sequential. If the access pattern is sequential, the storage system transfers contiguous tracks from the disks to cache ahead of time to improve cache hit rate. Due to this advance track transfer, the storage system shows the number of tracks transferred from the disks to the cache slot at DASD/CACHE of the SEQUENTIAL in TRANSFER OPERATIONS field in the storage system counters report, even though the access mode is not sequential.

- **IDCAMS LISTDATA STATUS:** The LISTDATA STATUS command generates status information for a specific device within the storage system. The storage systems report two storage sizes:
 - **Subsystem storage:** This field shows capacity in bytes of cache. For a storage system with more than one SSID, the cache is shared among the SSIDs instead of being logically divided. This strategy ensures backup battery power for all cache in the storage system. For the Hitachi storage systems, this field shows 75% of the total cache size.
 - **Nonvolatile storage:** This field shows capacity in bytes of random access cache with a backup battery power source. For the Hitachi storage systems, this field shows 25% of the total cache size.
 - **IDCAMS SETCACHE:** The storage systems support the IDCAMS SETCACHE commands, which manage caching for subsystem storage through the use of one command (except for REINITIALIZE). The following SETCACHE commands work for the subsystem storage across multiple SSIDs:
 - **SETCACHE SUBSYSTEM ON|OFF**
 - **SETCACHE CACHEFASTWRITE ON|OFF**
 - **SETCACHE NVS ON|OFF**
 - **SETCACHE DESTAGE**
-  **Note:** The SETCACHE REINITIALIZE command reinitializes only the logical subsystem specified by the SSID. You must issue the REINITIALIZE command once for each defined SSID.
- **DEVSERV PATHS:** The DEVSERV PATHS command is defined as the number of LVIs that can be specified by an operator (from 1 through 99). To display an entire storage system, enter the DEVSERV command for several LVIs, as follows:
 - **DEVSERV PATHS ,100 ,64**
 - **DEVSERV PATHS ,140 ,64**
 - **DEVSERV PATHS ,180 ,64**
 - **DEVSERV PATHS ,1C0 ,64**

z/VM® (VM/ESA) cache operations

When a storage system is managed under z/VM®, the following SET CACHE commands are effective across multiple SSIDs:

- **SET CACHE SUBSYSTEM ON|OFF**
- **SET NVS SUBSYSTEM ON|OFF**
- **SET CACHEFW SUBSYSTEM ON|OFF**
- **DESTAGE SUBSYSTEM**

z/VSE® (VSE/ESA) cache operations

When a storage system is managed under z/VSE®, the following CACHE commands are effective across multiple SSIDs:

- **CACHE SUBSYS=cuu, ON | OFF | STATUS**
- **CACHE SUBSYS=cuu, FAST, ON | OFF**
- **CACHE SUBSYS=cuu, NVS, ON | OFF**
- **CACHE SUBSYS=cuu, REINIT**



Note: SIMs indicating a drive failure may not be reported to the z/VSE® console (reference IBM® document GA32-0253). Since the RAID technology and dynamic spare drives ensure nonstop processing, a drive failure may not be noticed by the console operator. Since all SIMs are also logged on the SVP, you can use Device Manager - Storage Navigator to monitor the SIMs.

Chapter 4: zLinux operations

This chapter describes storage system operations in a Linux host environment.

Overview of zLinux operations

The storage systems support attachment to the following mainframe Linux operating systems:

- Red Hat® Linux for S/390 and zSeries
- SUSE Linux Enterprise Server for IBM® zSeries

For information about supported versions, Linux kernel levels, and other details, contact your Hitachi Vantara representative.

For details about FICON®/zHPF connectivity, FICON®/Open intermix configurations, and supported adapters, switches, and directors (for example, Brocade,) for the storage systems, please contact your Hitachi Vantara representative.

For information about preparing the storage system for Linux host attachment, please refer to the following IBM® publication:

- *Implementing Linux with IBM Disk Storage, SG24-6261*

The following links (current at time of document publication) will enable the reader to understand the details about the different methods for using Linux on mainframe hosts:

- ESS library
- Linux on S/390® hosts:
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/sg244987.html?Open>
- FCP use on zSeries hosts:
<http://www.redbooks.ibm.com/abstracts/sg246694.html?Open>

After physical installation of the storage system has been completed, you can configure the Linux system for operations with assistance as needed from the Hitachi Vantara representative. For specific information and instructions on configuring the storage system disk devices for Linux operations, please refer to the *Open-Systems Host Attachment Guide*.

Please contact your Hitachi Vantara account team for the latest information about platform and operating system version support.

Running Linux on S/390® or zSeries hosts

There are three methods to run Linux on an S/390® or zSeries host:

- Run Linux in “Native” mode on the machine with no other operating system.
- Create Logical Partitions (LPARs). Logically partition the hardware into a maximum of 15 separate LPARs per logical channel subsystem. You can, for example, run S/390® applications in one partition, z/VM® and z/VSE® applications in another partition, and Linux applications in a third partition.
- Use z/VM® guest support. Run Linux as a virtual machine using z/VM® to provide virtualization of central processing units, I/O subsystems, and memory.

It is possible to run hundreds of Linux systems on a single S/390® or zSeries host. For example, you can create a complete Linux server environment for each application developer and host production system all on the same S/390® or zSeries host.

For SUSE SLES or Red Hat Enterprise Linux, the storage systems support 31- and 64-bit Linux for zSeries on storage units as well as FICON®, ESCON® (USP V/VM only), and fibre-channel protocol (FCP) attached host systems.



Note:

zSeries FCP host systems running SUSE SLES or Red Hat Enterprise Linux can be attached through a switched-fabric configuration or a direct-attached configuration. The support for FCP direct attachment depends on the zSeries processor in use. IBM® Z9 processors and later support direct connection with FCP.

For details, please refer to the latest version of the following IBM® manual:

IBM System z Connectivity Handbook (SG24-5444), section 5.1.1 FICON modes and topologies

Attaching FC adapters on zSeries hosts running Linux

Linux solutions are available for the 31- and 64-bit environments. The availability of this option depends on the zSeries model and the Linux distribution. Fibre-channel support is available for both direct and switched attachment.

- **Linux for S/390 (31-bit):** Linux for S/390® is a 31-bit version of Linux. It also runs on zSeries models in 31-bit mode. The 31-bit limitation limits the addressable main storage to 2 GB.
- **Linux for zSeries (64-bit):** Linux on zSeries supports the 64-bit architecture on all zSeries models. The 64-bit support eliminates the 31-bit storage limitation of 2 GB.

FCP support

Linux for zSeries can access FCP and FICON[®]/zHPF controllers and devices with the appropriate I/O driver support. Linux can be run in a logical partition or as a guest operating system under z/VM[®] Version 4 Release 4 or later. For the most current list of supported operating system releases, see the Interoperability Matrix at <http://www.ibm.com/servers/storage/disk/ds8000/interop.html>.

A zSeries FCP channel is required to attach storage through a fibre-channel port. An FCP channel requires the same FICON[®] channel hardware that is used for FICON[®] channel. FICON[®] channel hardware can be designated as either a FICON[®] or FCP channel by specifying the correct channel path type (FC or FCP) in the IOCP channel path ID (CHPID) statement. The FICON[®]/zHPF requires a CHPID type FC and uses the FCP link transport. Defining FCP type channels in the IOCP reduces the total number of FICON[®] type channels that you can configure.

Accessing storage units using FC protocol from a zSeries host

1. Gather the prerequisite information about the SAN setup. See [Setting up storage units for zSeries hosts running Linux \(on page 75\)](#) for a list of the information that you must collect.
2. Set up the storage unit and configure the LUNs for use with Linux for zSeries. See [Setting up storage units for zSeries hosts running Linux \(on page 75\)](#) for the steps you must perform to set up the storage unit
3. Manually load the required modules and parameters to access the disks. See [Setting up a Linux system to use FC protocol devices on zSeries hosts \(on page 76\)](#) for the steps you must perform to view LUNs and set the Linux system to use FCP devices.
4. Make the disks permanently accessible. See [Adding permanent devices for zSeries hosts running Linux \(on page 77\)](#) for the steps you must perform to add permanent devices.

Setting up storage units for zSeries hosts running Linux

To set up the LUNs for use with your Linux system, you first must define the system to the storage unit. You can do this through the IBM[®] TotalStorage DS Storage Manager. The following instructions assume that you are familiar with the IBM[®] TotalStorage DS Storage Manager.

Before you begin

You must begin by first collecting the following software configuration information to prepare a Linux system for accessing the storage unit through a fibre channel:

- Host name of the server hosting the Linux system
- Device address and CHPID of the FCP port that is attached to the Linux machine. FICON[®]/zHPF requires a CHPID type FC.
- Worldwide port name (WWPN) of the FCP port on the zSeries
- WWPN of the fibre-channel port on the storage unit
- Fibre-channel port on the storage unit

This information can be obtained from the hardware management console (HMC), the IBM® TotalStorage DS Storage Manager, and the SAN switch, as shown in the following example:

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Procedure

1. Log on to IBM® TotalStorage DS Storage Manager, and in the navigation select **Real-time manager > Manage hardware > Host systems**.
2. From the **Select Action** list, select **Create**, then click **Go**.
3. From the **General host information** panel, complete the following fields for each fibre-channel host adapter. When you are finished, click **OK**.
 - Host type
 - Nickname
 - Description
4. From the **Define host ports** panel, specify the host ports for this host. Click **Add** to add each host port to the defined host ports table.
5. From the **Define host WWPN** panel, specify the worldwide port names for the selected hosts. When finished, click **Next**.
6. From the **Select storage images** panel, specify the storage image for attachment. Click **Add** to add the storage image to the selected storage images table. When finished, click **Next**.
7. From the **Specify storage image parameters** panel, specify the following parameters for each storage image.
 - Host attachment identifier v Volume Group for host attachment
 - Type of I/O port (any valid storage image I/O port or a specific I/O port)
8. Click **Apply** assignment button to apply the current attachment assignment. Use this button to go through this page for each host attachment identifier that you want to assign to the storage image. When you are finished, click **OK**.
9. From the **Verification** panel, verify the data that represents the new host attachments. When finished, click **Finish**.

Setting up a Linux system to use FC protocol devices on zSeries hosts

Before you begin

You must collect the following software configuration information to prepare a Linux system to access the storage unit through a fibre channel:

- Host name of the server hosting the Linux system
- Device address (and CHPID) of the FCP port that is attached to the Linux machine

- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit
- WWPN of the fibre-channel port on the storage unit

You can obtain this information using the hardware management console (HMC), the IBM® TotalStorage DS Storage Manager, and the SAN switch. The following example shows the prerequisite information that you must obtain to run FCP Linux on the zSeries.

```
The device number of the device that is defined in the IOCP for the FCP
channel on the zSeries
The SCSI ID starting at 1
The WWPN of the storage unit fibre-channel port
The SCSI LUN within Linux starting at 0
The FCP LUN of the target volume on the storage unit2
```

Procedure

1. Confirm that you have installed the following modules (at a minimum) to use the FCP device with the storage unit:
 - **qdio**: This is the same module is used for other qdio devices
 - **scsi_mod**: Defines SCSI core parameters
 - **zfcplib**: Provides FCP support for zSeries Linux
 - **sd_mod**: SCSI disk support1
2. Load the modules in the order shown in step 1. Use the **modprobe** command to load all modules. Except for the **zfcplib** module, you can load all modules without parameters.
3. The **zfcplib** module requires parameters to map the FCP devices on the storage unit. Each FCP requires the following parameters:

```
The device number of the device that is defined in the IOCP for the
FCP channel on the zSeries
The SCSI ID starting at 1
The WWPN of the storage unit fibre-channel port
The SCSI LUN within Linux starting at 0
The FCP LUN of the target volume on the storage unit2
```

4. Use the following command string to enter the mapping parameters to the module:

```
map="devno SCSI_ID:WWPN SCSI_LUN:FCP_LUN" map="0x0600
1:0x5005076300c89589 0:0x5301000000000000"
```

Adding permanent devices for zSeries hosts running Linux

Before you begin

Collect the following software configuration information to prepare a Linux system to access the storage unit through a fibre channel. This information can be obtained using the hardware management console (HMC), the IBM® TotalStorage DS Storage Manager, and the SAN switch.

- Host name of the server that is hosting the Linux system
- Device address and CHPID of the FCP port attached to the Linux machine
- Worldwide port name (WWPN) of the FCP port on the zSeries
- Fibre-channel port on the storage unit
- WWPN of the fibre-channel port on the storage unit

The following example shows the prerequisite information that must be obtained to run FCP Linux on the zSeries.

```
Linux hostname linuxa
CHPID of the FCP port on the z800 15
WWPN of the FCP port on the z800 50:05:07:64:01:40:01:7d
FC port on the storage unit Enclosure 3 Slot 1
WWPN of the FC port on the storage unit 50:05:07:63:00:c8:95:89
```

Procedure

1. Choose one of the following methods to add devices: write a script or manually add the device.
 - To add more than one device to your SCSI configuration, write a small script that includes all the parameters included. This is an optional step.
 - Alternatively, you can also add SCSI devices to an existing configuration with the **add_map** command. After using this command, you must manually map the devices in the SCSI stack.
2. To make the devices available permanently, after a reboot, you must create a new initial ramdisk (RAM that is accessed like a disk) that contains the necessary modules and parameter information. First, save the module parameters in the configuration file `/etc/zfcp.conf`, as shown in the following example.


```
The device number of the device that is defined in the IOCP for the
FCP channel on the zSeries
The SCSI ID starting at 1
The WWPN of the storage unit fibre-channel port
The SCSI LUN within Linux starting at 0
The FCP LUN of the target volume on the storage unit2
```

- a. Create a new ramdisk with the `mk_initrd` utility, and then run the `zipl` utility to update the initial program load record to point to the new ramdisk.
3. Enable multipath support using the logical volume manager kernel for Linux z/FCP.
 - a. Create physical volumes using the following commands: `pvcreate /dev/sdal` `pvcreate /dev/sdbl...`

- b. Put the physical volumes into a volume group using the following command:
`vgcreate fcpvg /dev/sda1 /dev/sdb1...`
 - c. Create as many logical volumes as you need using the following command:
`lvcreate --size 16G fcpvg`
 - d. Enable the alternate paths to the physical volumes using the **pvpath** command: `pvpath --path0 --enable y /dev/sda1 pvpath --path1 --enable y /dev/sda1`. If both paths are set to a weight of 0, they will load balance.
4. These configurations yield the SCSI device `/dev/sda - /dev/sdc`. This device is accessible on the first path and the SCSI device `/dev/sdd - /dev/sdf` accessed on the second path.

Chapter 5: FICON[®] and ESCON[®] migration

This chapter describes the process required to migrate a zSeries between FICON[®] and ESCON[®] adapters and between a FICON[®] bridge to a FICON[®] adapter.

 **Note:** FICON[®] support consists of hardware enhancements for enterprise servers, host software upgrades, DS8000 LMC, and adapters. The FICON[®] bridge performs the ESCON[®]-to-FICON[®] conversion.

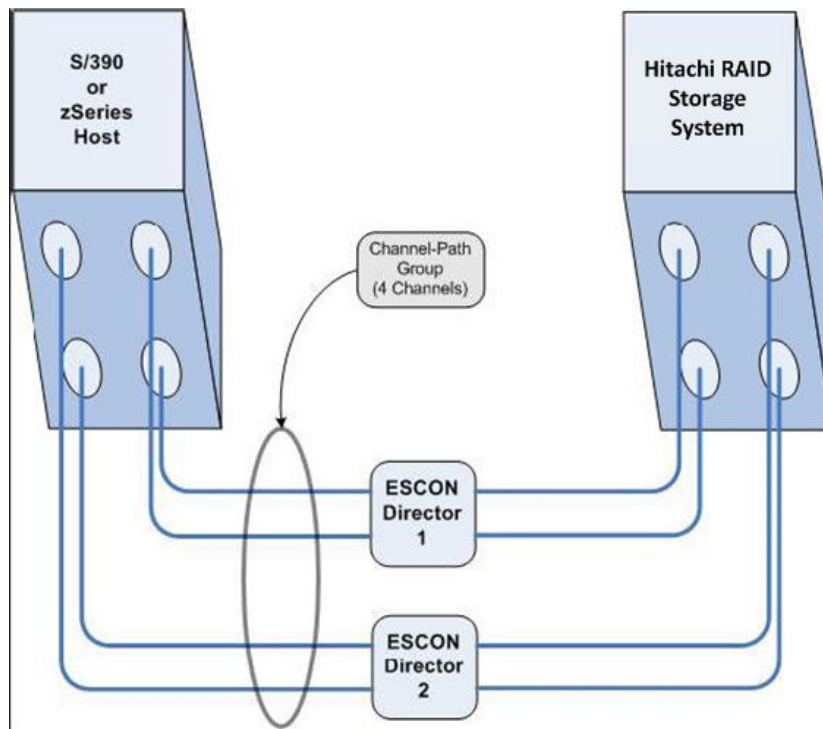
Migrating from ESCON[®] to native FICON[®]

This section covers the following topics:

- [Native ESCON configuration \(on page 80\)](#)
- [ESCON to FICON migration example for a zSeries or S/390 host \(on page 81\)](#)
- [FICON configuration example for a zSeries or S/390 host \(on page 82\)](#)

Native ESCON[®] configuration

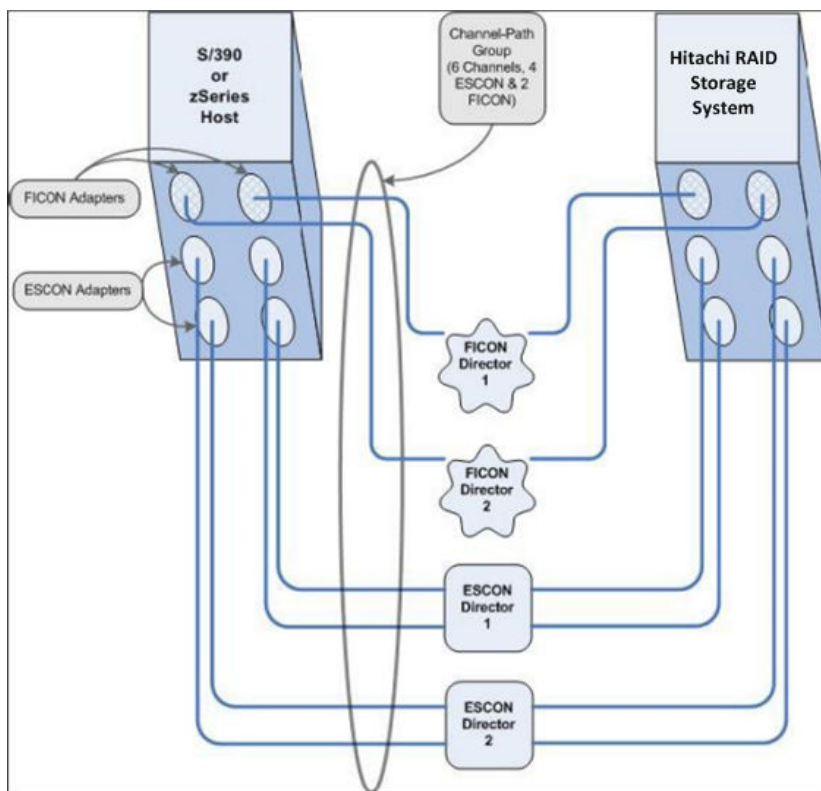
The following figure shows an example of a native ESCON[®] configuration. The configuration shows an S/390[®] or zSeries[®] host that has four ESCON[®] channels attached to the storage unit through two ESCON[®] directors.



The channels are grouped into a channel-path group for multi-pathing capability to the storage unit ESCON® adapters.

ESCON® to FICON® migration example for a zSeries or S/390 host

The following figure shows another example of an S/390® or zSeries host system with four ESCON® channels. In this example, two FICON® channels are added to an S/390® or zSeries host.



This figure shows four ESCON® adapters and two ESCON® directors. The illustration also shows the channel path group and FICON® directors through which the two FICON® adapters are installed in the storage unit. The two FICON® directors are not required.

You can improve reliability by eliminating a single point of failure. The single point of failure might be present if both FICON® channels are connected to a single FICON® director. You can connect the FICON® channels to the storage unit FICON® adapters directly, without directors.

This configuration gives you the most flexibility for future I/O changes. Because the channel path group has ESCON® and FICON® channel paths, it makes migrating easier. This intermixing of types of channel paths enables you to nondisruptively add FICON® channels to the host and to add FICON® adapters to the storage unit.



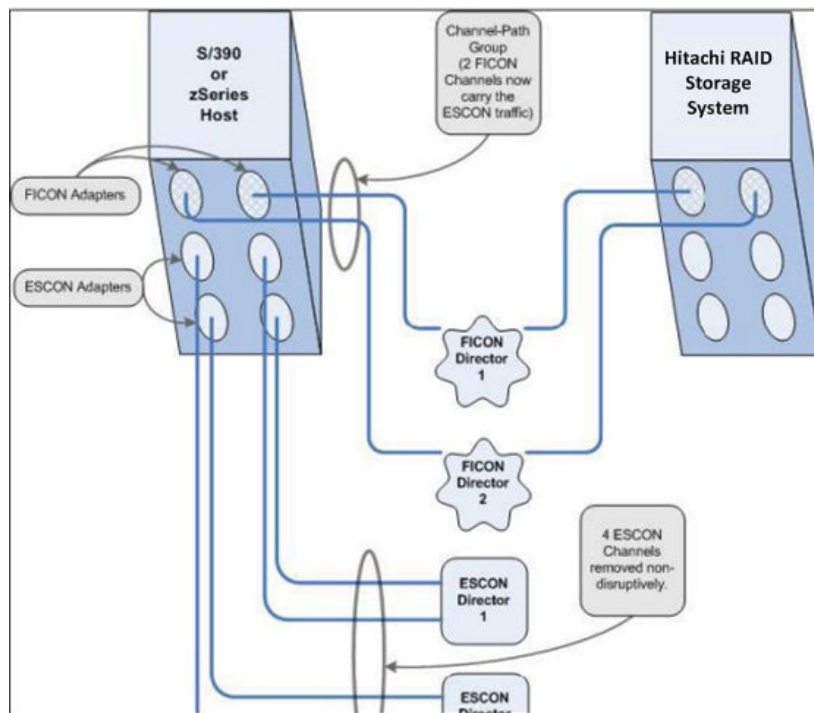
Note:

- This configuration is supported for migration only.
- Do not use this configuration for an extended period.
- Migrate from a mixed channel-path group configuration to an all FICON® channel-path group configuration.

FICON® configuration example for a zSeries or S/390® host

The following figure provides a FICON® configuration example for a zSeries or S/390® host that illustrates how to remove the ESCON® paths.

The S/390® or zSeries host has four ESCON® channels connected to two ESCON® directors. The S/390® or zSeries host system also has two FICON® channels.



You can remove the ESCON® adapters nondisruptively from the storage unit while I/O continues on the FICON® paths. By changing the channel-path group definition to include only the FICON® director paths you can complete the migration to the storage unit with two FICON® adapters. You can also retain the ESCON® channels on the S/390® or zSeries host system so you can access other ESCON® controllers. You also can keep the ESCON® adapters installed in the storage unit to connect to other S/390® or zSeries hosts.

Migrating from a FICON® bridge to a native FICON® attachment

This section covers the following topics:

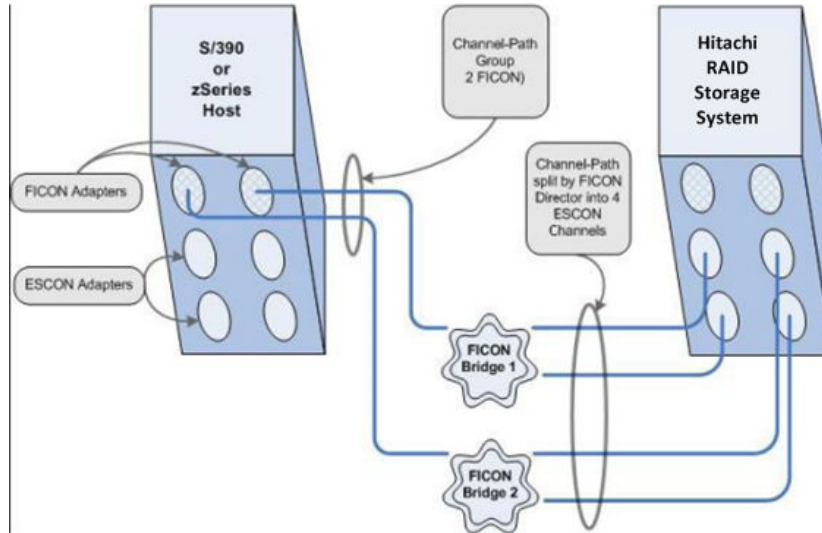
- [FICON bridge overview \(on page 83\)](#)
- [FICON bridge configuration \(on page 84\)](#)
- [FICON and FICON channel mixed configuration \(on page 84\)](#)
- [Native FICON configuration on a zSeries or S/390 host \(on page 85\)](#)
- [Registered state-change notifications \(RSCNs\) on zSeries hosts \(on page 85\)](#)

FICON® bridge overview

The FICON® bridge is a feature card of the ESCON® Director 9032 Model 5. The FICON® bridge supports an external FICON® attachment and connects internally to a maximum of eight ESCON® links. The volume on these ESCON® links is multiplexed on the FICON® link. You can perform the conversion between ESCON® and FICON® on the FICON® bridge.


FICON® bridge configuration

The following figure shows an example of a FICON® bridge. It consists of an S/390 or zSeries host with two FICON® channels attached to two FICON® bridges. You can attach the storage unit through the channel-path group to four ESCON® links.



FICON® and FICON® channel mixed configuration

The following figure provides an example of an S/390® or zSeries host system with one FICON® channel and one FICON® director through a channel-path group and FICON® host adapter to the storage unit. The example also shows that the S/390® or zSeries host system with one FICON® channel and one ESCON® director with a FICON® bridge through a channel-path group and two ESCON® adapters.



In the example, one FICON® bridge was removed from the configuration. The FICON® channel that was connected to that bridge is reconnected to the new FICON® director using the storage unit FICON® adapter. The channel-path group was changed to include the new FICON® path. The channel-path group is now a mixed ESCON® and FICON® path group. I/O operations continue to the storage unit across this mixed path group. Access to the storage unit devices is not interrupted because all the actions are nondisruptive.

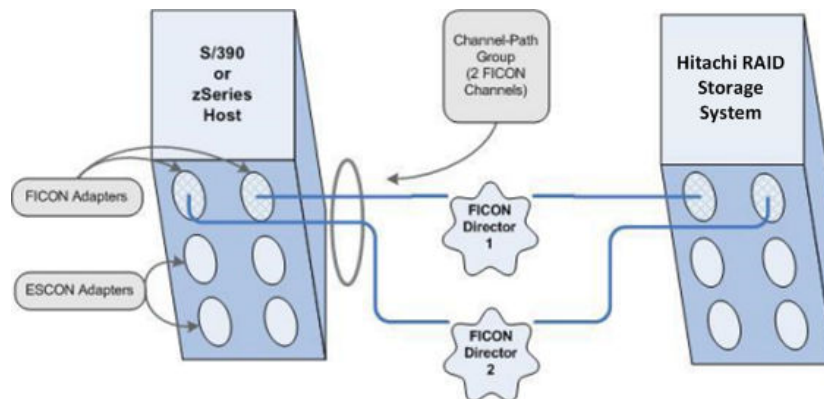


Note:

- This configuration is supported for migration only.
- Do not use this configuration for an extended period.
- Migrate from a mixed channel-path group configuration to an all FICON® channel-path group configuration.

Native FICON® configuration on a zSeries or S/390® host

The following figure shows an S/390® or zSeries host system with two FICON® channels connected to two FICON® directors through a channel-path group to two FICON® adapters. Note that the second bridge has been removed and a FICON® director has been added. The channel-path group has only the FICON® paths.



Registered state-change notifications (RSCNs) on zSeries hosts

McDATA and CNT switches ship without any configured zoning. This un-zoned configuration enables the default zone on some McDATA switches. This configuration enables all ports in the switch with a fibre-channel connection to communicate with each other and to receive registered state change notifications about each other. You can set the zones.

Here are some recommendations:

- If you have FICON®-only environments, do not disable the default zone.
- If you have FICON® environments and plan to add (or already have) fibre-channel intermix on the switch, ensure that all FICON® host adapters and all FICON®-capable device ports are in an explicitly defined common zone.

Appendix A: Notes on using zDAC with VSP G1x00 and VSP F1500

For information and instructions on using the zDAC function, see the IBM® Hardware Configuration Definition User's Guide. The following limitations and restrictions apply when the zDAC function is used for the Hitachi Virtual Storage Platform G1x00 and Hitachi Virtual Storage Platform F1500 storage systems.

This appendix covers the following topics:

- [Notes on executing the zDAC function \(on page 86\)](#)
- [Message output \(on page 86\)](#)
- [Specifications of open volumes \(on page 87\)](#)

Notes on executing the zDAC function

The zDAC function has the following specifications. Consider the following when executing the zDAC function:

- The zDAC function in the direct connection configuration supported with z/OS® V2R1 is not supported.
- Volumes such as 3390-V and volumes that are set as system disks that cannot be placed online from mainframe hosts are not defined.
- When FABRIC DISCOVERY is performed for the already created configuration, IFCC or MIH may occur depending on the usage. Note that new CUADD cannot be added to the created configuration.
- When "Force full mode discovery" of "Discovery and Autoconfiguration Options" is set to "NO", CU in which UA=x00-x10 and xFF are not installed may not be discovered. In this case, set "Force full mode discovery" to "YES".
- Whether Alias VOL is installed or not depends on the emulation type of Base VOL when Alias is set. If Base VOL is not installed, Alias VOL is also not installed. For example, with 2107 emulation, if Base VOL is 3390-A, both Base VOL and Alias VOL are not installed.

Message output

When the zDAC function is executed, the following message may be output in CU in which UA is not installed from x00. Note that the execution of zDAC is not affected.

```
CHP=chpid STAT=SUBCHANNEL PATH AVAILABLE, BUT DEVICE NOT OPERATIONAL
```

Specifications of open volumes

In IBM® storage, both mainframe and open volumes are defined in the configuration. However, for VSP G1x00 and VSP F1500, only volumes that can be placed online from the mainframe host can be defined in the configuration.

Hitachi Vantara Corporation



Corporate Headquarters

2845 Lafayette Street

Santa Clara, CA 95050-2639 USA

www.HitachiVantara.com | community.HitachiVantara.com

Regional Contact Information

Americas: +1 866 374 5822 or info@hitachivantara.com

Europe, Middle East, and Africa: +44 (0) 1753 618000 or info@emea@hitachivantara.com

Asia Pacific: + 852 3189 7900 or info.marketing.apac@hitachivantara.com