Draft Document for Review June 9, 2025 7:45 am SG24-8581-00

Redbooks
ibm.com/redbooks

IBM z17 Configuration Setup

Ewerson Palacio
Franco Pinto
Nelson Oliveira
Ryotaroh Sawada
Hongshou Liu
Martin Soellig
Patterson Monteiro
Elcio Pineschi
Volkmar Langer

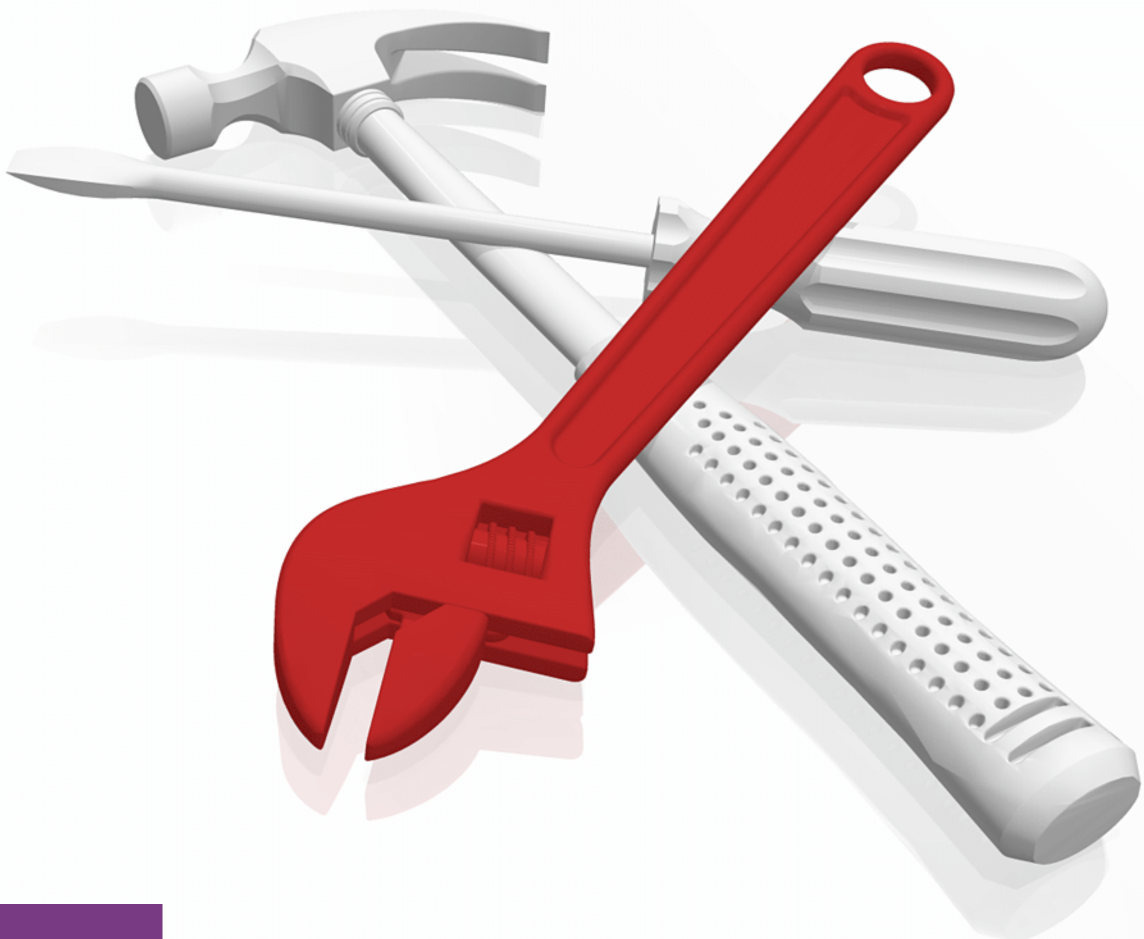

IBM Z

Redbooks

Note: Before using this information and the product it supports, read the information in “Notices” on page ix.

First Edition (June 2025)

Contents

Notices ix
Trademarks x

Preface xi
Now you can become a published author, too! xii
Comments welcome xii
Stay connected to IBM Redbooks xiii

Chapter 1. Introduction 1
1.1 High-level goal 1
1.2 Scope 2
1.3 Configuration tools 6
1.4 Preview of IBM z17 server changes 6

Chapter 2. Planning considerations 7
2.1 Scenario descriptions 8
2.1.1 Upgrading an existing IBM Z server to an IBM z17 8
2.1.2 Installing a new IBM z17 8
2.1.3 Planning for the scenarios 8
2.2 Key tools 9
2.2.1 IBM Resource Link 10
2.2.2 Hardware Configuration Definition 11
2.2.3 IBM Z Connectivity Mapping Tool 12
2.2.4 HCD and the CMT 13
2.3 Extra tools 13
2.3.1 Input/output configuration program 13
2.3.2 Worldwide Port Name Prediction Tool 14
2.3.3 Coupling Facility Structure Sizer 14
2.3.4 Power estimation tool 14
2.3.5 Shared Memory Communications Applicability Tool 15
2.3.6 IBM Z Batch Network Analyzer tool 15
2.4 Hardware Management Appliance 16
2.4.1 Activation profiles 16
2.4.2 Cryptographic configuration 17
2.4.3 LPAR group control 18
2.4.4 Consoles and terminals 18
2.4.5 Hardware Management Appliance 19
2.4.6 Hardware Management Console considerations 20
2.4.7 Support Element settings 24
2.4.8 Network Time Protocol 25
2.4.9 Precision Time Protocol 25
2.4.10 Server Time Protocol 26
2.4.11 Pulse per second 26
2.4.12 External Time Sources (ETS) and STP planning considerations 26
2.5 IODF configuration 29
2.5.1 Channel subsystems 29
2.5.2 Logical partitions 29
2.5.3 Storage connectivity 32
2.5.4 Network connectivity 34

2.5.5 Coupling and timing links 37
2.5.6 Planning considerations for hardware data compression 38

Chapter 3. Preparing for a new IBM z17 system 39
3.1 Supported hardware features 40
3.2 Saving and restoring Open Systems Adapter-Express configuration data 41
3.3 Upgrading an IBM z16 to an IBM z17 while maintaining the serial number 41
3.3.1 Scenario overview. 42
3.3.2 Creating the work IODF from the current 3931 production IODF 43
3.3.3 Repeating the 3931 processor to be replaced 43
3.3.4 Coupling Link information messages 45
3.3.5 Deleting any unsupported items in the repeated 3931 46
3.3.6 Changing the 3931 to a 9175 and deleting the 3931 54
3.3.7 Deleting the 3931 processor definition 55
3.3.8 Reconnecting the CF channel paths that were not migrated 57
3.3.9 Additional steps and tasks 57
3.4 Installing a new IBM z17 scenario overview 57
3.4.1 Creating a work IODF from the current production IODF 58
3.4.2 Adding the 9175 processor 59
3.4.3 Additional steps and tasks 62

Chapter 4. Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool 63
4.1 Creating a copy of an existing IODF from an IBM z16 server. 64
4.1.1 Change the Processor Type 67
4.1.2 Validating the work input/output definition file. 68
4.2 Creating the IOCP file for use with the CMT 70
4.3 Assigning CHIDs to CHPIDs by using the CMT 73
4.4 Importing the CFR file and the IOCP input file into CMT. 74
4.5 Building the initial Overview 75
4.6 Assigning missing feature types 78
4.7 Resetting Incompatible (Hardware - I/O) Entries. 80
4.7.1 Assigning PCHID entries to FUNCTION IDs (FID) 81
4.8 Creating CMT reports 82
4.8.1 CHPID Report 84
4.9 Creating an updated IOCP file 87
4.10 Additional steps and processes 88

Chapter 5. Building the production input/output definition file and setting up the central processor complex 89
5.1 Building the new production IODF 90
5.2 Writing the input/output configuration program to the old CPC by using HCD 93
5.3 Creating a reset profile on the Support Element 101
5.3.1 Background activities that occurred 101
5.3.2 Building the reset profile and pointing it to the required IOCDS 101
5.3.3 Setting up and verifying the reset profile. 104
5.4 Creating an image profile on the Support Element 105
5.4.1 Image Profile: General page 105
5.4.2 Image Profile: Processor page 106
5.4.3 Image Profile: Security page 107
5.4.4 Image Profile: Storage page 108
5.4.5 Image Profile: Options page 109
5.4.6 Image Profile: Load page 110
5.4.7 Image Profile: Crypto page 111
5.4.8 Image Profile: Time Offset 112

5.4.9 Image Profile: Saving . . . 112
5.5 Performing a Power on Reset on the new CPC . . . 113
5.5.1 Coupling Facility Links . . . 113
5.5.2 Remote dynamic activation of I/O configurations for Stand Alone Coupling Facilities (SACFs), Linux on IBM Z, and z/TPF . . . 117
5.6 Server Time Protocol configuration . . . 122
5.7 Building and verifying load (IPL) profiles . . . 124
5.8 Building and verifying LOADxx members in SYS#.IPLPARM . . . 127
5.9 Communicating information about the new CPC . . . 128

Chapter 6. Configuring network features . . . 129
6.1 Configuring OSA for OSA-Express and Network Express Adapter . . . 129
6.1.1 Open Card Specific Advanced Facilities . . . 131
6.1.2 Display or alter MAC address . . . 132
6.1.3 Enable or disable ports . . . 133
6.1.4 Set card mode . . . 133
6.2 Verifying OSA configuration . . . 133
6.3 Configuring Shared Memory Communications (SMC-D and SMC-R) . . . 134
6.3.1 Defining Shared Memory Communications over RoCE (SMC-R) . . . 135
6.3.2 Defining Shared Memory Communications Direct (SMC-D) . . . 135

Chapter 7. Defining console communication . . . 137
7.1 Preparing a console definition . . . 137
7.2 Defining the OSA-ICC . . . 138
7.3 Configuring a new OSA-ICC configuration by using OSA Advanced Facilities . . . 138
7.3.1 Saving and restoring the OSA-ICC configuration . . . 145
7.4 Verifying the OSA-ICC definition . . . 151
7.4.1 z/OS commands . . . 151
7.4.2 OSA-ICC console initial window . . . 152

Chapter 8. Preparing for IBM Parallel Sysplex and Server Time Protocol . . . 153
8.1 Preparing for Parallel Sysplex . . . 154
8.2 Preparing for non-sysplex system time synchronization . . . 154
8.3 Server Time Protocol overview . . . 154
8.3.1 External Time Source . . . 156
8.3.2 Configuring the HMC as an NTP server . . . 157
8.4 HMC 2.17.0 (Manage System Time task) . . . 158
8.4.1 Setting CTN member restrictions . . . 161
8.4.2 HMC operations to add a CPC to the CTN . . . 161
8.4.3 Verifying the new CTN configuration . . . 163

Chapter 9. Defining Coupling Facility links . . . 165
9.1 Coupling connectivity options for Parallel Sysplex on IBM z17 . . . 166
9.1.1 Preparing to define Coupling Facility links . . . 166
9.2 Defining Coupling Express3 Long Reach . . . 168
9.2.1 CE3 LR 10Gb (CL5) connections: . . . 169
9.2.2 CE3 LR 10Gb (CL5) verifying the configuration: . . . 170
9.2.3 CE3 LR 25Gb (CL6) connections: . . . 174
9.2.4 CE3 LR 25Gb (CL6) verifying the configuration: . . . 176
9.3 Defining Integrated Coupling Adapter Short Reach . . . 178
9.3.1 ICA SR: Verifying the configuration . . . 180
9.4 Defining an STP timing-only link by using ICA SR . . . 183
9.4.1 STP timing-only link: Implementation . . . 183
9.4.2 STP timing-only links: Verifying the configuration . . . 186

9.5 CF LPAR setup and Coupling Facility Control Code Level 26 186
9.5.1 Coupling Facility Control Code Level 26 187
9.6 Dynamic I/O for Stand-alone Coupling Facility 187
9.6.1 IBM Processor Resource/System Manager solution. 188

Chapter 10. Specialized features 189
10.1 Crypto Express8S 190
10.1.1 Crypto Express8S overview 190
10.1.2 Planning for a Crypto Express8S configuration 193
10.1.3 Configuring Crypto Express8S 195
10.1.4 Handling cryptographic coprocessors by using ICSF 217
10.2 Virtual Flash Memory 223
10.2.1 VFM overview 224
10.2.2 Planning for VFM configuration. 224
10.2.3 Configuring VFM 224
10.2.4 VFM management 226

Chapter 11. Adding logical partitions and operating system configurations 237
11.1 Preliminary actions 238
11.2 Defining a new OSCONFIG 238
11.3 Defining a new LPAR 239
11.3.1 Define LPAR using HCD 239
11.3.2 Define image profile on HMA 241

Chapter 12. Adding storage devices 243
12.1 Defining more I/O by using HCD 244
12.1.1 Defining FICON switches (directors, storage area networks, and storage area network switches) 244
12.2 FICON CHPIDs, switches, and DASD CUs 246
12.2.1 Defining FICON CHPIDs and connecting them to a FICON switch 246
12.2.2 Defining FICON CHPIDs for a direct connection to a 2107 control unit 249
12.2.3 Defining FICON CHPIDs for switch connections to a 2107 control unit 252
12.2.4 Defining 3390B devices to an OSCONFIG and Eligible Device Tables and Esoteric device groups 258

Chapter 13. Adding network devices 265
13.1 Summary of network devices defined in this chapter 266
13.2 OSC definitions (Integrated Console Controller). 266
13.2.1 Defining OSC CHPIDs 266
13.2.2 Defining OSC CHPID connections to an OSC control unit 269
13.2.3 Defining 3270-X devices to an OSCONFIG 270
13.2.4 Defining 3270-X devices to the Nucleus Initialization Program (NIP) Console List within an OSCONFIG 273
13.2.5 Verifying OSC definition 277
13.3 OSD definitions (Queued Direct I/O - QDIO) 277
13.3.1 Defining OSD CHPIDs 277
13.3.2 Defining OSD CHPID connections to an OSA control unit 281
13.3.3 Defining OSA devices to an OSCONFIG(OSD CHIPDs) 283
13.3.4 Verifying the OSD definition 286
13.4 Network Express for channel type OSH 287
13.4.1 Defining OSH CHIPDs 287
13.4.2 Defining OSH CHPID connections to an OSA control unit 291
13.4.3 Defining OSA devices to an OSCONFIG(OSH CHIPDs) 291
13.4.4 Verifying the OSH definition 291

13.5 IQD CHPIDs for HiperSockets 292
13.5.1 Defining IQD CHPIDs 293
13.5.2 Defining IQD CHPID connections to an IQD control unit 296
13.5.3 Defining IQD devices to an OSCONFIG 298
13.5.4 Verifying the IQD definition 300

Chapter 14. Adding coupling connectivity 303
14.1 Defining more I/O by using an HCD 304
14.2 Coupling Facility logical partitions, CS5, CL5, CL6, and ICP CHPIDs 304
14.2.1 Defining Coupling Facility LPARs in a channel subsystem. 304
14.2.2 Defining CS5 CHPIDs. 305
14.2.3 Defining a Coupling Facility link with CS5 CHPIDs. 309
14.2.4 Defining CL5 and CL6 CHPIDs. 312
14.2.5 Defining a Coupling Facility link with CL5 or CL6 CHPIDs 318
14.2.6 Defining ICP CHPIDs 322
14.2.7 Defining a Coupling Facility link with ICP CHPIDs 326

Chapter 15. Adding Peripheral Component Interconnect Express devices 331
15.1 Defining PCIe functions by using HCD 332
15.2 PCIe feature definitions. 332
15.2.1 Overview 332
15.2.2 Defining an ISM PCIe function 334
15.2.3 Defining a NETH PCIe function. 339
15.2.4 Defining a zHyperLink PCIe function 345

Chapter 16. Defining Channel to Channel Connections 353
16.1 Understanding FCTC 353
16.2 Defining FCTC 354
16.3 Verifying FCTC 357
16.3.1 Checking the status using z/OS commands 357
16.3.2 Checking the status by using the SE: 357

Appendix A. Additional material 359
Locating the web materials 359
Using the web materials. 359
Downloading and extracting the web material 359

Abbreviations and acronyms 363

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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

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

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

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

The performance data and customer examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

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

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

Trademarks

Preface

This IBM® Redbooks® publication helps you install, configure, and maintain the IBM z17™ (machine type 9175) systems. The IBM z17 systems offers new functions that require a comprehensive understanding of the available configuration options. This book presents configuration setup scenarios, and describes implementation examples in detail.

This publication is intended for systems engineers, hardware planners, and anyone who needs to understand IBM Z® configuration and implementation. Readers should be familiar with IBM Z technology and terminology. For more information about the functions of the IBM z17 systems, see *IBM z17 Technical Introduction*, SG24-8580 and *IBM z17 (9175) Technical Guide*, SG24-8951.

This book was produced by a team of specialists from around the world working at IBM Redbooks, Poughkeepsie Center.

Ewerson Palacio is an IBM Redbooks Project Leader. He holds Bachelors's degree in Math and Computer Science. Ewerson worked for IBM Brazil for over 40 years and retired in 2017 as an IBM Distinguished Engineer. Ewerson co-authored a large number of IBM Z publications, created and presented at IBM Redbooks seminars around the globe.

Franco Pinto is a Z Client Architect with IBM Switzerland. He has over 30 years of experience in the mainframe and IBM z/OS fields. His areas of expertise include IBM Z technical pre-sales covering mainframe sizing, installation planning, and providing support on existing and new IBM Z functions. In his current role he supports clients and business partners how to best make use of the wide portfolio of IBM products and services.

Nelson Oliveira is a Consultant IT Specialist with IBM Brazil. He has 33 years of experience in the field of mainframe technology. His areas of expertise include IBM z/OS, Job Entry Subsystem 2 (JES2), IBM Parallel Sysplex, High Availability, IBM Geographically Dispersed Parallel Sysplex (IBM GDPS), and IBM Z platform.

Ryotaroh Sawada is a Consultant IT Specialist in Japan. He has 17 years of experience in technical support for IBM Z clients. His areas of expertise include IBM Z hardware, z/OS, and Systems Management on the mainframe. He has co-authored several IBM Z Redbooks publications, from the IBM z14 ZR1 to the IBM z16.

Hongshuo Liu is an IT Specialist in IBM Japan. She has 9 years of experience in the IBM Z field. Her areas of expertise include IBM Z hardware and IBM z/OS. She is a member of a team that supports IBM Z hardware and software for multiple major banks.

Martin Soellig is an IBM Z Technical Specialist in Germany. He has 34 years of experience working in the IBM Z field. He holds a degree in mathematics from the University of Hamburg. His areas of expertise include IBM z/OS and IBM Z Servers™ hardware, specifically in Parallel Sysplex® and GDPS® environments, and also in cryptography on IBM Z.

Patterson Marletta Monteiro is a Consultant Client Availability Leader in Brazil. He has 23 years of experience working with banks in Brazil on IBM Z platform. He holds a degree in Computer Science from UNISA and a MBA in Communication and Networks from University of São Paulo (Escola Politécnica). His areas of expertise include IBM z/OS and Communications server.

Elcio Pineschi is an IT Project Manager at Caixa Economica Federal, one of the major banks in Brazil. He holds a degree in Eletrical Engineering from Universidade Estadual do Rio de

Janeiro and a MSC degree in Computer Engineering from Universidade Federal do Rio de Janeiro. He has 25 years of experience working with IBM mainframes system support.

Volkmar Langer is an IBM MVS™ Senior Systems Programmer at Dataport in Germany. He has more than 35 years of experience in z/OS and IBM Z hardware. For over 25 years, he has been responsible for the concepts and design of Dataport's z/OS Sysplex environments, including LPAR design and disaster recovery concepts. His areas of expertise include z/OS, Parallel Sysplex, JES2, HCD, HMC and Software Management (SMP/E).

Thanks to the following people for their contributions to this project:

Robert Haimowitz, Bill White
IBM Redbooks, Poughkeepsie Center

Patty Driever, David Hollar, Dave Surman
IBM Poughkeepsie

Jason Stapels, John Troy, Robert Miekle
IBM US

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

http://www.redbooks.ibm.com/rss.html

1

Introduction

This chapter describes the high-level goal of this book. It covers scenarios that were devised based on best practices. The scenarios are described in subsequent chapters along with the tools that are used to implement the configurations.

> **Notes:** The IBM z17™ generation is available in one configuration:
>
> - The IBM z17 Model ME1 is built in a 19-inch form factor designed to fit inside a standard 19-inch rack that scales from 1 to 4 frames, depending on the configuration. IBM z17 ME1 ensures continuity and upgradeability from IBM z16 A01 and IBM z15 T01. It has five orderable features: Max43, Max90, Max136, Max183, and Max208.

This chapter includes the following topics:

- High-level goal
- Scope
- Configuration tools
- Preview of IBM z17 server changes

1.1 High-level goal

The goal of this book is to help you plan for and complete the configuration tasks for a successful installation of IBM z17 systems, machine type 9175. It covers the planning and preparation tasks that are needed from when an IBM z17 system is delivered and physically installed, up to the point when a logical partition (LPAR) is ready to be activated.

This book describes the planning considerations and configuration examples in detail from Hardware Management Appliance (HMA), Support Element Appliance (SE), and input/output definition file (IODF) perspectives.

1.2 Scope

Before you perform the planning and preparation tasks that are covered in this book, some activities must be completed:

- Customers Configuration Design: Together with your team, IBM provides design and configuration information for the installation of the IBM z17 system that you plan to purchase.
- IBM Order to Manufacturing: Your IBM representative orders the approved configuration. IBM makes the machine configuration available for download as a CFReport. The CFReport file can be obtained from the IBM Resource Link® website (you must authenticate by using your registered IBMid) and by using a Configuration Control Number (CCN) provided by your IBM representative.
- Physical installation: With support from IBM, the new order machine or a frame roll upgrade to an IBM z17 system is physically installed.

> **Note:** A frame roll upgrade includes the exchange of the existing frames with newly delivered frames and CPC drawers. Although the machine serial number is retained.

> **Note:** The IBM z17 server does not support HMCs as separate physical equipment. The HMC functions are provided by ordering the HMA feature. The optional HMA feature provides redundant HMC functions, and is supported by the SE server in the IBM Z server A frame. IBM recommends that at least one HMA feature is ordered for a maximum of two per site.

- Trusted Key Entry (TKE) installation: With support from IBM, the (optional) TKE workstations are installed. If necessary, contents such as user profiles and API settings are migrated (if you replace the TKEs).

Figure 1-1 shows the steps that are required for each distinct scenario when preparing for the installation of the IBM z17 system, which includes the following steps:

- Upgrading an existing IBM z15 or z16 system to an IBM z17 system
- Installing a new IBM z17 system

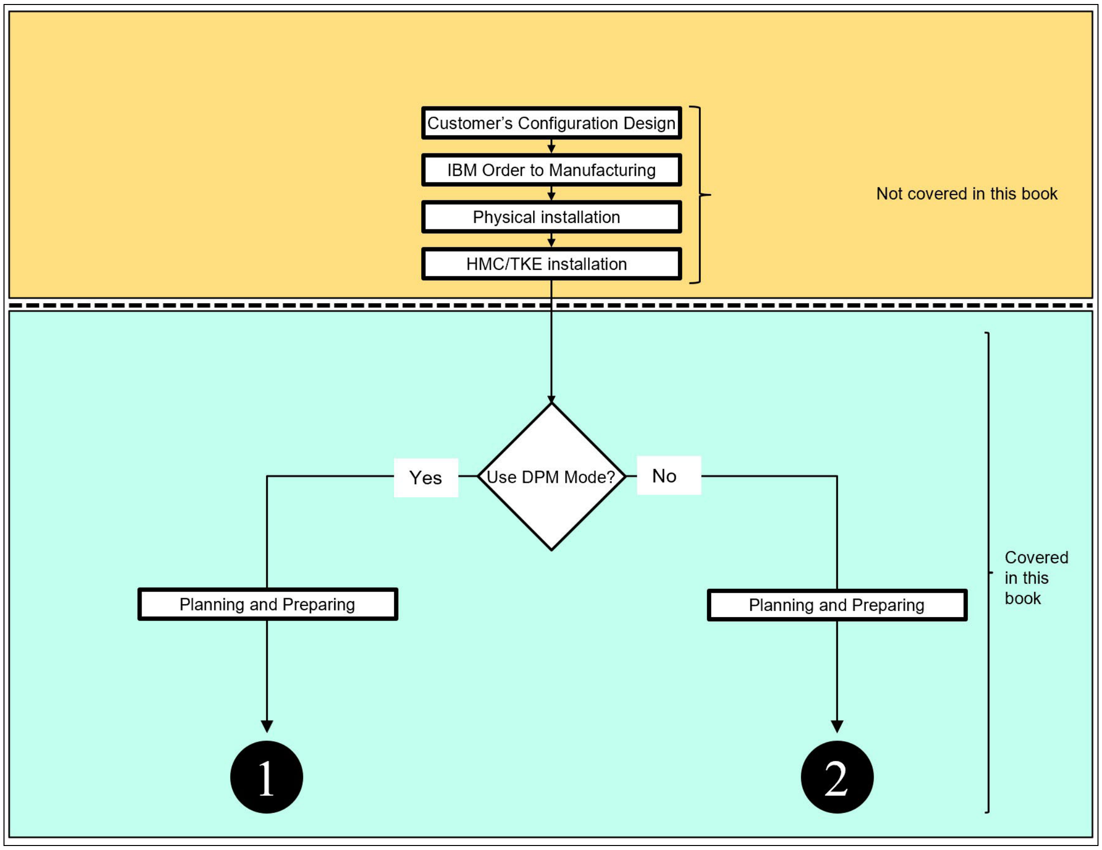

Figure 1-1 Showing the topics that are covered in this book

The flowchart in Figure 1-1 is divided into two different main sets of task streams:

- The upper stream (marked in light yellow) is based on actions that should be performed before hardware arrival, they are not covered in this book.
- The lower stream (marked in light green) also accounts for the usage of Dynamic Partition Manager (DPM) as an option.

With DPM, system administrators have a quicker and simpler way to deploy Linux on IBM Z, IBM z/VM®, Kernel-based Virtual Machine (KVM), and Secure Service Container (SSC) LPARs. DPM is a wizard-like configuration method that runs in the HMA.

Important: When DPM is enabled, the IBM z17 system cannot run z/OS, 21st CS VSE^{n} R6.3.1, and z/TPF LPARs.

The flowchart in Figure 1-2 shows the required tasks to install DPM and define LPAR operating characteristics by using the HMC function. The flowchart is divided into two task streams:

- The stream marked in light yellow in the flowchart is based on actions that must be performed by IBM on the SE before the IBM z17 system is handed over to you.
- The stream marked in light green describes the configuration flow for a partition by using the DPM application. With the input that is provided to DPM, a configuration is activated that is used on the IBM z17 system to host an OS.

The actions that are defined in the two streams must be performed in sequence.

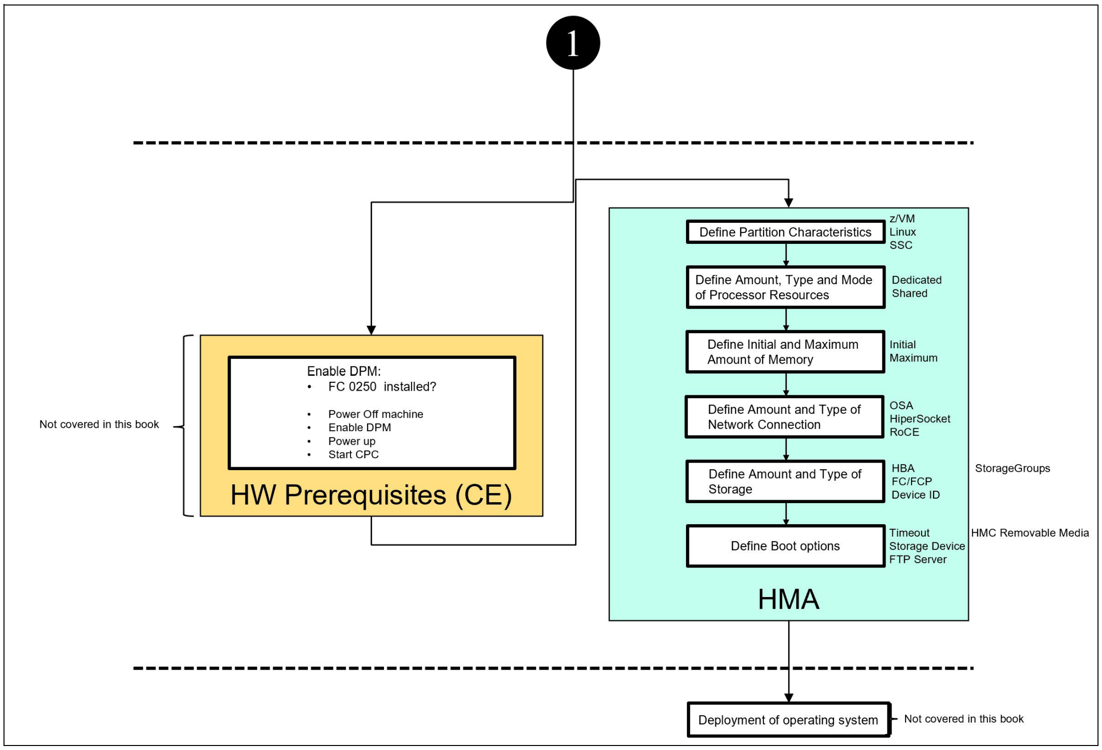

Figure 1-2 Installation flowchart that is applicable to both an IBM z17 system upgrade and a new installation that uses DPM

DPM automatically discovers and displays the system resources that are available for use in your Linux on IBM Z, z/VM, KVM, and SSC LPARs. When using DPM, the partition configuration data is created, which contains a description of all I/O functions and features that are used on the IBM z17 system, all compute and memory resources, and all crypto assignments.

> **Note:** This book does not cover scenarios that use DPM. For more information about the usage of DPM, see *IBM Dynamic Partition Manager (DPM) Guide*, SB10-7182.

The flowchart that is presented in Figure 1-3 on page 5 describes more tasks that must be done to complete the installation. The flowchart is divided in two different task streams:

- The left side of the flowchart is based on actions that must be performed on the HMC function or the SE function of the HMA.
- The right side of the flowchart is based on definitions in the hardware configuration management program.

The actions that are defined in the two streams can have dependencies between them.

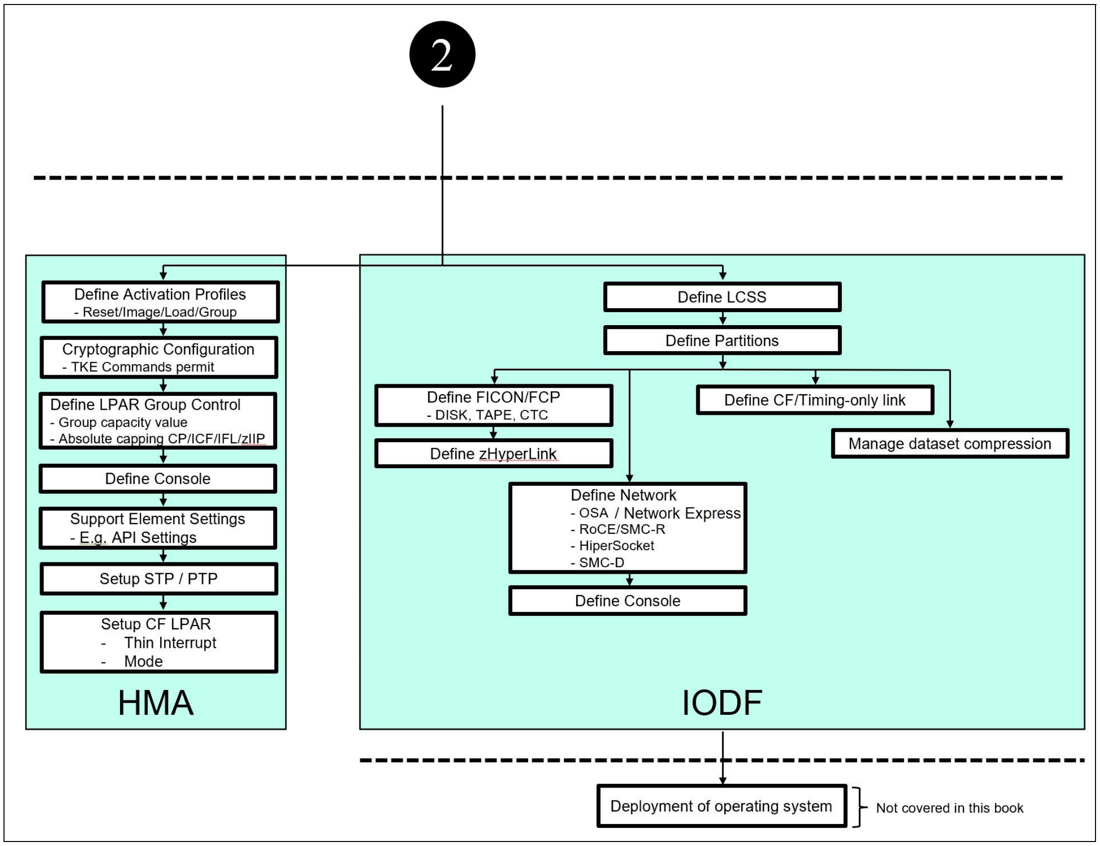

Figure 1-3 Installation flowchart for an IBM z17 system (applicable to a new system or an upgrade)

HMA communicates with the SE server (physically installed in the A frame of the IBM z17 system), which provides communication with the IBM z17 hardware. On the HMC function, you must set some parameters so that you can activate the number of LPARs that run a supported OS. To create an IODF (see Figure 1-3), you must perform a set of activities in an application (such as Hardware Configuration Definition (HCD), which needs a running z/OS system). The IODF can be created on a different system than the target system. There are multiple HMC/SE and IODF tasks that must be planned and prepared for. For more information, see *I/O Configuration Using z/OS HCD and HCM*, SG24-7804.

These flowcharts are intended to act as a checklist rather than a step-by-step procedure. The steps in this book should provide enough information for you to replicate the approach in your environment.

For more information about how to deploy an OS (z/OS in this case), see *Mainframe from Scratch: Hardware Configuration and z/OS Build*, SG24-8329.

1.3 Configuration tools

Several tools are provided by IBM that can help you achieve a successful IBM z17 installation. Whenever possible throughout this book, lists are provided to help you go through the steps that are required to complete a specified task.

In addition to the tools and lists that are provided in this document, ensure that the planning and configuration steps align with other technical departments within your organization, such as storage and network administration, and with the capacity (workload) planning and cryptographic and security teams.

Configuration tools, like HCD, Connectivity Mapping Tool (CMT), and the HMA appliance, are covered in Chapter 2, “Planning considerations” on page 7.

1.4 Preview of IBM z17 server changes

The IBM z17 server is the next generation of CPC for z/OS and other IBM Z operating systems. It has been announced in client letter AD25-0015. Compared to previous hardware generations (IBM z15 server to IBM z16 server) IBM z17 implemented changes in the I/O connectivity design of the server hardware, mainly based on the new Data Processing Unit (DPU) for I/O Acceleration. The implementation of the DPU is transparent to the user of this book.

New I/O adapter cards have become available with the IBM z17 server. The integration of the new features into the IODF will be described in this book. We will describe approaches on how to upgrade an I/O configuration from an existing z15 or z16 processor to a z17 as well as the steps needed to make a new build IBM z17 server available.

We will illustrate the changes in the configuration tasks needed particularly in the networking area. With the IBM z17, new FICON Adapter Cards will be available. The specifics and the configuration tasks will be described in the following chapters. Carry forward of OSA Express adapters from previous IBM Z server generations is possible. While this eases the upgrade process, the switch to Network Express adapters that are new with IBM z17 server provide a number of advantages. The Network Express adapters are simpler to define, and they provide the opportunity to consolidate to fewer ports, thus reducing the number of adapters.

2

Planning considerations

This chapter describes planning and configuration considerations for the IBM z17 system. Whenever possible, worksheets that support the planning tasks are provided. Throughout this book, we provide various definition examples by using Hardware Configuration Definition (HCD) as the preferred method for the I/O configuration. Other tools, such as Hardware Configuration Manager (HCM) and input/output configuration program (ICP IOCP), are mentioned for reference only.

Naming: The IBM z17 ME1 system is targeted by this publication. Throughout this chapter, we might refer to this machine as IBM z17.

This chapter also provides a short overview of tools that IBM provides to help with configuring your IBM z17 system, and information about where to obtain the tools and their intended use.

This chapter includes the following topics:

- Scenario descriptions
- Key tools
- Extra tools
- Hardware Management Appliance
- IODF configuration

2.1 Scenario descriptions

Throughout this book, we use two distinct scenarios to explain the tasks and procedures that are involved to successfully install and configure an IBM z17 system:

- Upgrading an existing IBM Z server to an IBM z17
- Installing a new IBM z17

2.1.1 Upgrading an existing IBM Z server to an IBM z17

This scenario assumes that an existing IBM Z is upgraded by using a miscellaneous equipment specification (MES) to a IBM z17 system. The scenario includes a planned outage period during the physical upgrade of the machine. The software environments that are supported by the machine type prior to the upgrade are not available during the upgrade period. The serial number of the machine remains the same after the upgrade to the IBM z17 system.

2.1.2 Installing a new IBM z17

This scenario assumes that a new IBM z17 system is installed in an existing mainframe environment. The IBM z17 system is physically installed along an existing IBM Z machine. After the installation of the IBM z17 system successfully completes and the system is handed over by IBM, the software environment on the machine to be replaced must be stopped, and recabling actions must be performed.

When recabling is complete, post-installation activities must be performed, and the software environment can be brought back online on the new (IBM z17) system. An outage can be planned for this scenario, and a new serial number must be considered, so software keys for the new system must be available before the migration.

2.1.3 Planning for the scenarios

In the first scenario, the physical platform identity (machine serial number) that is configured remains the same. No hardware configuration files must be physically migrated to another platform. No changes to the software licenses are required for products that are tied to the machine serial number.

Note: Software licensing might change depending on the machine capacity.

In the second scenario, the physical platform that is configured changes. Hardware configuration files must be prepared on the existing machine, and must be migrated to the new IBM z17 system together with the attached cabling. The serial number changes with the activation of the IBM z17 system, which means that planning and preparing for software license changes must be considered beforehand.

In both scenarios, we assume that bringing up the existing features and functions has the highest priority. Adding new features and functions that are acquired with the system upgrade or installed in the new IBM z17 system have a lower priority. The elapsed time of the planned outage can vary, depending on the approach that is chosen in either scenario.

In both scenarios, some information must be obtained before starting the process of changing to or installing the new IBM z17 system:

- The new processor ID: The processor ID is used to assign a unique name to identify the processor in the HCD. For more information, see *z/OS HCD Users Guide*, SC34-2669.
- The CFReport file: The CFReport file is downloadable from IBM Resource Link by entering a Configuration Control Number (CCN). The CCN is provided by your IBM representative.
- The system serial number: If a new IBM z17 system will be installed, a new serial number is provided by your IBM representative.

Also, IBM does not provide fiber optic cables as features on the IBM z17 system. Therefore, a complete analysis of the I/O connectors that are used on existing systems that are upgraded to an IBM z17 system must be made to ensure that the appropriate fiber optic cabling is installed.

An equivalent study should be part of your preparation to install a new IBM z17 system so that all cabling is delivered to the data center before the installation date.

All required cables for the IBM z17 should be identified and placed on order. Labeling all cables is required for the installation. At a minimum, the labels should identify the physical channel ID (PCHID) number.

If you already received the configuration and PCHID reports from IBM, define your coupling links to fit your planned configuration to your new or upgraded central processor complex (CPC).

Note: Physical Hardware Management Consoles (HMCs) are not orderable or supported for a new build IBM z17 or machines upgraded to z17. The HMC functions for an IBM z17 is provided by the IBM Z Hardware Management Appliance (HMA) optional feature (Feature Code 0355).

2.2 Key tools

IBM provides several tools to help with the complexity of configuring an IBM Z platform. This section summarizes the various tools that are available, and briefly outlines their benefits for the planning process.

Table 2-1 lists the machine types for the IBM Z platform. The examples in this book use tools, such as the HCD and IBM Z Connectivity Mapping Tool (CMT), which refer to the machine type instead of names. For more information, see Chapter 4, “Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool” on page 63.

Table 2-1 Machine types for IBM Z platforms

Name	Machine type
IBM z17 ME1	9175
IBM z16 A02 and IBM z16 AGZ	3932

Name	Machine type
IBM z16 A01	3931
IBM z15 T02	8562
IBM z15 T01	8561

2.2.1 IBM Resource Link

The first step in planning for the installation of the IBM z17 is to access IBM Resource Link. Access to the IBM Resource Link can be requested on the IBM Resource Link site. Your IBM representative can assist you with the registration process. After you register for an IBM ID, you can customize your profile to accommodate the platforms that you are responsible for. After completing the registration process with IBM, you will be provided with credentials to logon to your site's part of IBM Resource Link.

On the Resource Link website, you have access to various resources and tools that are designed to help the installation process. Several tools are available to simplify the installation process of an IBM z17 system. Even if you worked with most of these tools before, be sure to check for the latest versions that are relevant to the IBM z17.

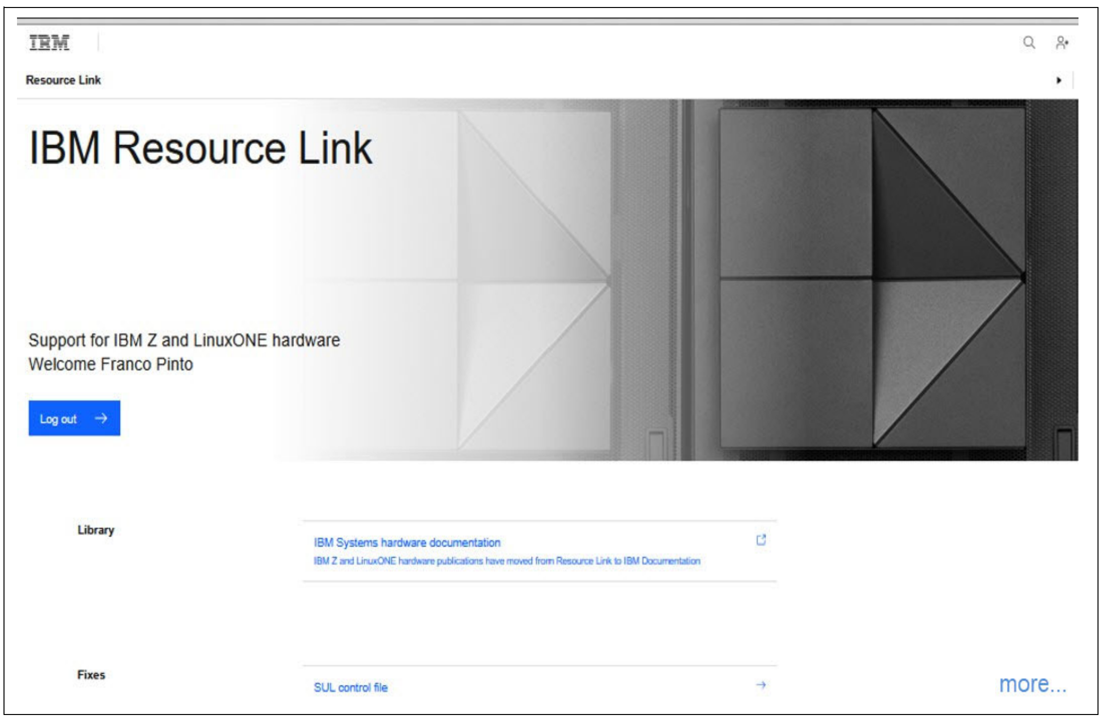

Figure 2-1 shows the screenshot of the Resource Link main page

- The **Library** part of the page links you to the hardware documentation in IBM docs
- The **Fixes** part points you to where you can download SUL Control Files, Firmware downloads and can request Remote Code downloads
- The **Problem solving** links you to the IBM Z and LinuxONE Security Portal, as well a to the Call Home Connect Cloud webpage

- The **System/Solution Assurance** points you to the Technical Delivery Assurance (TDA) forms
- The **Tools** part links you to the support tools such as:

  1090 support (z/PDT)
  IBM Z Connectivity Mapping Tool
  Customized installation manual search
  Customized planning aids search
  File upload
  Host key document search
  IBM Z Application Programming Interfaces
  Power and weight estimation
  WWPN tool

- The **Capacity on Demand** part points you to the Customer Initiated Upgrade part of IBM Resource Link.

IBM Call Home Connect Cloud is IBM's web application that allows IBM hardware clients to view and monitor key status indicators about their Call Home-enabled IBM hardware assets. Offered at no additional cost to all IBM hardware clients, the website displays information about:

- Critical cases and alerts
- Warranty and maintenance contract status
- Last contact status
- Current software and firmware levels and upgrade recommendations
- Asset details, such as:
  - Summary data about open and closed cases, with links to IBM Support to view them
  - Detailed alerts about various Call Home-related events
  - Related assets
  - Current installed code levels and code level installation history (for selected products)
  - Links to product-specific support information
  - System Z machine information reports

2.2.2 Hardware Configuration Definition

HCD is a component that runs on IBM z/OS and IBM z/VM. It supplies an interactive dialog to generate the input/output definition file (IODF) and the input/output configuration data set (IOCDS).

Consider using HCD or HCM to generate the I/O configuration rather than writing your own IOCP statements. HCD performs validation as you enter the data, thus minimizing the risk of errors. This book provides examples for using HCD, with some examples that used HCM (see “Hardware Configuration Manager” on page 12).

Note: New hardware (an IBM z17 system) requires program temporary fixes (PTFs) to enable definition support Unit Information Modules (UIM's) in HCD.

When defining devices in HCD, the hardware features can be selected according to the physical setup of the devices that are attached to the IBM z17. Detailed forms and charts that describe the environment facilitate the planning process.

For more information about HCD, see IBM Documentation.

Hardware Configuration Manager

HCM provides a GUI to HCD and the associated IODF. HCM runs on a workstation and can also define and store more information about the physical hardware to which the IODF is defined.

HCM does not replace HCD. It is used with HCD and the associated IODF. However, HCM can be used in a stand-alone mode after an IODF is built and the configuration files (`IODF##.HCM or IODF##.HCR`) are created on your HCM workstation.

For more information about HCM, see z/OS and z/VM HCM User's Guide.

2.2.3 IBM Z Connectivity Mapping Tool

The IBM Z Connectivity Mapping Tool (CMT) provides a mechanism to map CHPIDs to PCHIDs on an IBM z17 system. Using the CMT is preferable to manually mapping the CHPIDs to PCHIDs. Use of the CMT provides the best availability practices for a configuration.

Two files are needed to obtain an IODF file containing the correct PCHID numbers by using the CMT:

- A copy of a production IODF file (also called a work IODF file) without PCHID numbers. For more information about how to obtain this file, see Chapter 4, "Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool" on page 63.
- The CFReport file reflecting the physical configuration of the ordered IBM z17, which is obtained from the Resource Link website. To obtain the CFReport, the corresponding machine CCN is required. The CCN is generated by your IBM Client Representative when building the order for your configuration.

2.2.4 HCD and the CMT

The HCD process flow for a new IBM z17 installation is shown in Figure 2-2.

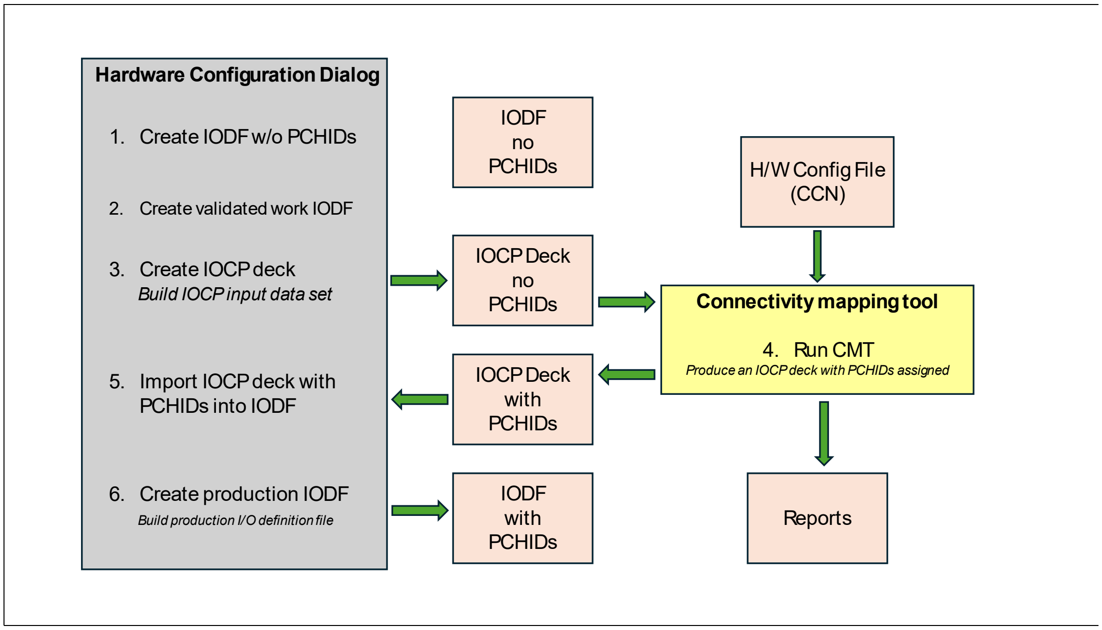

Figure 2-2 CMT: I/O configuration definition flow for a new installation

Part of the actions that are described in Figure 2-2 might also be valid for an upgrade, depending on the hardware configuration of the upgraded machine.

To download the CMT, log in to the Resource Link website with a registered IBMid and select **Tools**.

For additional information please refer to: *Connectivity Mapping Tool Users Guide*, GC28-7058 available in IBM Resource Link.

For more information about how to use the CMT, see Chapter 4, “Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool” on page 63.

2.3 Extra tools

The additional tools that are described in this section are not used in this book. However, they can help speed up the process of planning and configuring specific features or functions outside of this book.

2.3.1 Input/output configuration program

ICP IOCP Version 7 Release 1.0 or later is required for an IBM z17.[1] You can define the IBM z17 configuration by using only IOCP. However, using HCD is a best practice because of its verification and validation capabilities. By using ICP IOCP, it is possible to write an IOCDS in preparation for a CPC upgrade.

[1] ICP IOCP Version 6 Release 1 Level 2 (6.1.2) is required for the IBM z16 A02 and IBM z16 AGZ.

For more information about the changes and requirements for ICP IOCP, see *Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7183.

2.3.2 Worldwide Port Name Prediction Tool

The Worldwide Port Name Prediction Tool for Fibre Channel Protocol (FCP) Channels helps prepare configuration files that are required or generated by the IBM Z platform when FCP Channels are configured. This tool helps during the installation of new systems and system upgrades.

One of the most important configuration parameters are worldwide port names (WWPNs), which uniquely identify physical or virtual Fibre Channel (FC) ports. They are typically used in storage area network (SAN) switches to assign the corresponding ports to zones of a SAN. They are used in storage subsystems to grant access from these ports to specific storage devices that are identified by logical unit numbers (LUNs).

The capability of the WWPN Prediction Tool is extended to calculate and show WWPNs for both virtual and physical ports before system installation.

The WWPN Prediction Tool, which applies to a CPC in PR/SM mode, is available for download from IBM Resource Link and applies to all Fibre Connection (IBM FICON®) channels that are defined as CHPID type FCP (for communication with SCSI devices) on an IBM z17. You can access the tool on IBM Resource Link by using your IBMid or going through the IBM Resource Link home page and then selecting **Tools ® WWPN Tool**.

WWPN Persistence

The FCP WWPNs are determined based on the I/O serial number of the CPC, the IOCDS configuration details (for N_Port ID Virtualization and WWPNs), and the PCHID values (for physical WWPNs). When Feature Code 0099[2] (WWPN Persistence) is ordered as part of a new or upgraded configuration for an IBM z17 system, the I/O serial number part of the WWPN for the new IBM z17 system is the same serial number as for the source machine configuration.

2.3.3 Coupling Facility Structure Sizer

Moving to a new IBM z17 system means migrating to a higher Coupling Facility Control Code (CFCC) level (level 26). If your existing Coupling Facility (CF) data structures are adequately sized and you want to know how much these structures might need to grow to accommodate the same workload at the new CFCC level, you can use the current structure sizes to calculate the new sizes. The Coupling Facility Structure Sizer (CFSizer) Tool helps you evaluate the sizing of the CF structures.

Use the CFSizer tool to plan more accurately the amount of storage that must be allocated for CF partitions. You can access the tool at the CFSizer page.

2.3.4 Power estimation tool

The power estimation tool is a web-based tool that you use to estimate the power consumption for your IBM Z platform. For the IBM z17, based on the features that are

[2] Feature Code 0099 is an information-only feature code that is required to initiate an ordering option to maintain IO serial numbers when a newly purchased system replaces an existing system within the customer's data center. This option eliminates the need to reconfigure zoning in SAN switches and LUN masking in storage controllers.

installed on the machine, the tool also estimates its weight, air flow, exhaust temperature, individual phase currents, and power cord wattage.

For access to this tool, see IBM Resource Link.

2.3.5 Shared Memory Communications Applicability Tool

You can use the Shared Memory Communications Applicability Tool (SMCAT) to evaluate TCP/IP network traffic and use the evaluation to plan for any potential use of SMC-R. To make the evaluation, you do not need to enable the SMC-R function on any system or enable RoCE Express (on an IBM z15 or a IBM z16 system) nor Network Express (on IBM z17). Shared Memory Communications Applicability Tool[3].

This tool is very important on IBM z15 and IBM z16 systems since to exploit SMC-R the RoCE Express adapter card needs to be acquired. On IBM z17 SMC-R communication is enabled by associating a Network Express adapter card with a NETH PCIe function. For more information, see the Shared Memory Communications Applicability Tool.

2.3.6 IBM Z Batch Network Analyzer tool

IBM Z Batch Network Analyzer (IBM zBNA) is a no charge, “as is”, workstation based productivity tool designed to understand the batch window and estimate the CPU upgrade effect by performing “what if” analyses. Use IBM zBNA to evaluate the impact of new technology exploitation such as zEDC, DFSMS Encryption, zHyperLink, and DFSORT IBM Z Sort. IBM zBNA provides powerful, graphic demonstration of the z/OS batch window, which can be automatically captured in a report file using the Named Favorites feature.

The zBNA tool and its Users Guide can be downloaded from the IBM Z Batch Network Analyzer (zBNA) Tool website.

[3] This link points to the z/OS V3R1 documentation site. The SMCAT is also available in z/OS V2R5, and V2R4 with support extension.

2.4 Hardware Management Appliance

This section introduces the configuration and management tasks that are available on the Service Element (SE) function as part of the Hardware Management Appliance (HMA) of the SE server.

Note: The IBM z17 server does not support HMCs as separate physical equipment. The HMC functions are provided by ordering the HMA feature. For more information about HMA details, see 2.4.5, “Hardware Management Appliance” on page 19.

2.4.1 Activation profiles

Activation profiles must be customized by using the HMA. Activation profiles are required for CPC and image activation. They are used to tailor the operation of a CPC and are stored in the SE appliance that is associated with the CPC. There are four types of activation profiles:

- Reset: A reset profile is used to activate a CPC and its images.
- Image: An image profile is used to activate an image of a previously activated CPC.
- Load: A load profile is used to load a previously activated image with a control program or operating system (OS).
- Group: A group profile is used to define the group capacity value for all logical partitions (LPARs) belonging to that group.

The default profiles of each of these types are provided. The *Activate* task activates the CPC or image. Initially, the *Default* profile is selected. You can specify an activation profile other than Default. This feature provides the capability to have multiple profiles, for example, one for every IOCDS file that is managed by the CPC.

Reset profile

Every CPC in the processor cluster needs a *reset profile* to determine the mode in which the CPC Licensed Internal Code (LIC) is loaded and how much physical memory is available. Using the reset profile, you must provide the order in which the LPARs are activated during a Power on Reset (POR). The maximum number of reset profiles for each CPC is 26.

For more information about how to define a reset profile, see 5.3, “Creating a reset profile on the Support Element” on page 101.

Image profile

Each LPAR has an image profile. The image profile determines the number of CPs that the image uses and whether these CPs are dedicated to the partition or shared. It can also assign the amount of initial storage and reserved storage that are used by each partition, and points to the IOCDS slot in the SE that has the I/O configuration to load in the HSA. Depending on the SE model and machine type, the maximum number of image profiles that are allowed for each CPC can be in the range 64 - 255.

The parameters for each LPAR define these settings:

- General: The Profile name and its description, the partition identifier, and the mode of operation
- Processor: The number of logical central processors (CPs), IBM Z Integrated Information Processors (zIIPs), and the initial processing weight that is assigned to the LPAR

- Security: The security options for this LPAR, the BCPii permissions, the counter facility security options, the sampling facility security options, and the CP Assist for Cryptographic Functions (CPACF) key management operations
- Storage: The total amount and the initial amount of memory that is assigned to this LPAR and the Virtual Flash Memory (VFM) allocation to this LPAR (if VFM is present)
- Options: The I/O priority, defined capacity options, and the CP management cluster name
- Load: The load type and address parameters that are necessary to run an IPL for this LPAR
- Crypto: The Crypto Express parameters (see 2.4.2, “Cryptographic configuration” on page 17)

> **Note:** To help you gather the necessary input, a worksheet is provided with this book. For more information about downloading the worksheet that is associated with this material, see Appendix A, “Additional material” on page 359.

For more information about how to define an image profile, see 5.4, “Creating an image profile on the Support Element” on page 105.

Load profile

A *load profile* is needed to define the channel address of the device from which the OS is loaded. Depending on the SE model and machine type, the maximum number of load profiles for each CPC is 64 - 255.

Group profile

A *group profile* defines the group capacity value that can be customized to determine the allocation and management of processor resources that are assigned to the LPAR in a group. This profile does not contain the names of the LPAR images that make up the group.

2.4.2 Cryptographic configuration

The activation profile that you use to activate an LPAR prepares it for running software products that can use the Crypto Express feature. Using the feature’s cryptographic facilities and functions requires customizing the LPAR’s activation profile to accomplish these tasks:

- Install the CPACF Data Encryption Standard Enablement feature (FC 3863) if you are planning to use Integrated Cryptographic Service Facility (ICSF).
- Give the LPAR access to at least one Crypto Express feature. This goal is accomplished by selecting from the Usage Domain Index and the Cryptographic Candidate list.
- Load the LPAR with an OS, such as z/OS, that supports using cryptographic functions.

For more information about the cryptographic features, see 10.1, “Crypto Express8S” on page 190.

2.4.3 LPAR group control

This task allows you to view or change a group assignment for logical partitions. It displays the group name, member partitions, and group capacity value that can be customized in determining the allocation and management of processor resources assigned to the group. It also allows changing a group assignment dynamically for active logical partitions.

Here are methods that can be used to limit the processor capacity usage for a single LPAR or a group of LPARs and help you control software cost:

- Edit Group Capacity: Use this window to view or change a group assignment for logical partitions. This window displays the group name, member partitions, group capacity value, and absolute capping setting that can be customized in determining the allocation and management of processor resources assigned to the group. The system can manage the group in such a way that the limit for Group Capacity in MSU per hour is not exceeded.
- Absolute Capping: Use this field to change the absolute capping of logical partitions in a group that share processors. The absolute capping can be none or a number of processors value from 0.01 to 255.0. Absolute Capping is managed by IBM Processor Resource/System Manager (PR/SM) and, is independent of the OS running in the capped LPARs.

Both methods (Group Capacity and Absolute Capping) can be used concurrently and with LPAR capping.

Consider reevaluating the parameters in a scenario where the values must be migrated from a previous generation of the IBM Z to an IBM z17 system.

> **Tip:** Capacity management by using capping technologies is an ongoing process that must be monitored and adjusted over time. Temporary or permanent capacity changes must also be considered when using capping technologies.

A good overview of the capping technologies and 4-hour rolling average (4HRA) optimization can be found in Capping Technologies and 4HRA Optimization.

2.4.4 Consoles and terminals

Consoles and terminals are provided by the OSA-Express Integrated Console Controller (OSA-ICC). OSA-ICC requires OSA Express7S 1.2 GbE SX/LX features to support TN3270 enhancements (TN3270E) and non-Systems Network Architecture (SNA) distributed function terminal (DFT) 3270 emulation. Planning for an IBM z17 OSA-ICC implementation requires input from several disciplines within an organization:

- IBM Z I/O subsystem configuration
- OS configuration
- OSA-Express feature configuration
- Ethernet local area network (LAN) configuration
- Client TN3270E configuration

Note: IBM z17 does not support OSA-Express 1000BaseT adapter, except as a Carry Forward from zIBM z15.
Channel type OSC must be defined on the OSA-Express7S GbE 1.2 SX or LX (FC0442 or FC0443) only.

IBM z17 supports the use of the Transport Layer Security protocol 1.2 (TLS 1.2) and 1.3 (TLS 1.3) to establish secure connections to terminals and consoles).

In HCD, the OSA-Express feature must be defined to operate as CHPID type OSC. The configuration requirements are as follows:

- IBM Z I/O subsystem configuration: The same basic rules for adding an OSA-ICC adapter apply as to any other new device.
- OS configuration: To have a Nucleus Initialization Program Console available, make sure that the correct device number is defined in the HCD OS “Work with consoles” dialog.

During an upgrade from an existing IBM Z platform to an IBM z17 system, the same OSA-ICC definitions can be used for the new machine as on the source configuration.

The following planning topics must be considered:

- Reserve at least one OSA-Express7S 1.2 GbE SX or LX port with the correct support to be defined as channel type OSC.
- Define 3270-X Devices in HCD to act as system consoles or terminals.
- Use OSA Advanced Facilities to configure the sessions.

The OSA-Express feature also requires configuration tasks to be performed on the HMC function by using the OSA Advanced Facilities task. Collect information for the following parameters before starting the configuration activities on port 0 and port 1:

- OSA-ICC server: Name, Host IP address, Mask, TCP port number, secure TCP port number, MTU size, Gateway IP address and minimum TLS version
- OSA-ICC session definitions: Channel subsystem (CSS), the Multiple Image Facility (MIF) (LPAR) ID, Device number, LU-name, client IP address, IP Filter, the session type, defer host disconnect (DHD), response mode (RSP), and read timeout (RTO)

Note: Consider defining multiple sessions per LPAR to allow access for several users concurrently. A maximum of 120 valid subchannels can be used on a OSC channel path.

For an upgrade of an existing IBM Z server to an IBM z17 system, these definitions can be exported from the source machine by using onboard HMC functions and imported back again after the upgrade is complete.

For more information about the definitions and implementation details, see Chapter 7, “Defining console communication” on page 137.

2.4.5 Hardware Management Appliance

The Hardware Management Appliance (HMA) is a closed system appliance, which means that no other applications can be installed on it. The HMA runs a set of management functions.

The HMC code runs as an application on two high performance top-of-rack servers. HMC functions are peers and support data replication. The HMA can be on IBM z17, z16 or z15. One HMA (Feature Code 0355) feature code provides two HMAs. The HMA was introduced with the IBM z15. They run on both of the two integrated 1U top-of-rack SE servers.

Figure 2-3 shows the physical location of the redundant 1U rack-mounted SE servers supporting the HMA on a fully equipped IBM z17 machine.

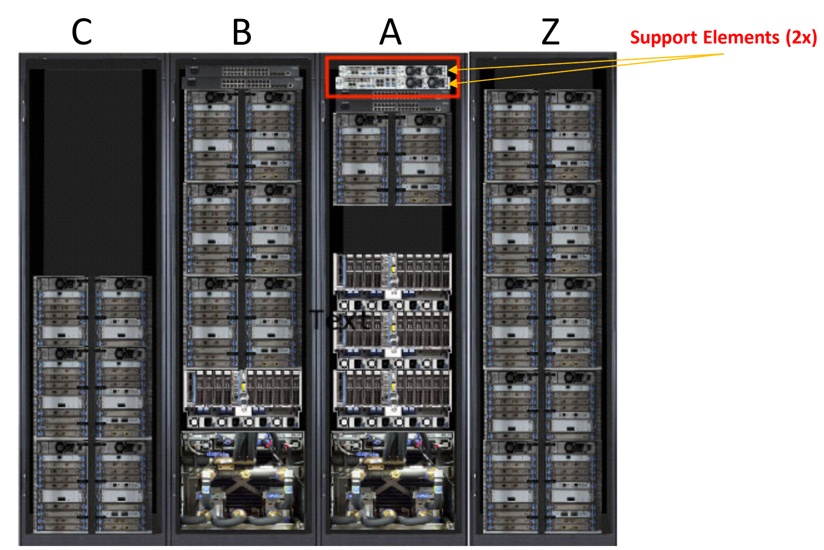

Figure 2-3 Back view of z17 server showing the location of the SE servers

2.4.6 Hardware Management Console considerations

With IBM z17 and HMA, the SE Console function is a licensed application that provides the tasks that you use to monitor and operate your system. The appliance is included with each SE server that is installed on the top of the IBM z17 A frame. One SE function runs as the designated Primary (or active) SE and the other as the designated Alternative (or backup) SE. As with the HMC function, the SEs are closed systems, and no other applications can be installed on the same environment.

Figure 2-4 on page 21 shows the evolution of the HMC/SE environment, including the HMA offering that was introduced with IBM z15.

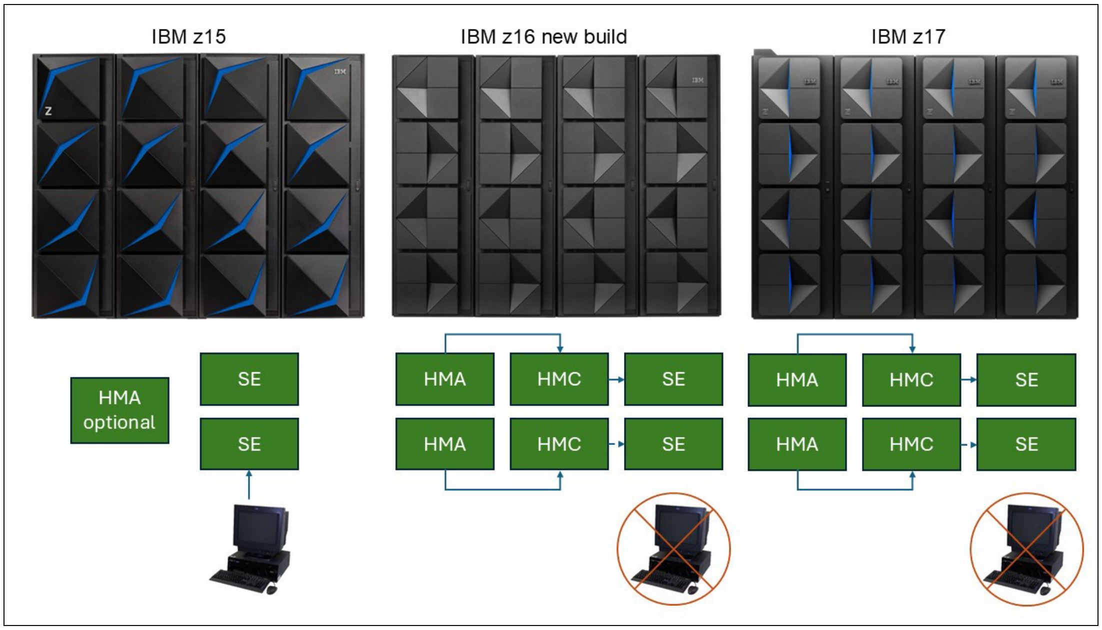

Figure 2-4 Evolution of the HMC/SE environment

The HMC function is used to set up, manage, monitor, and operate one or more CPCs. It manages IBM Z hardware, its LPARs, and provides support applications. At least one HMC function is required to operate an IBM Z. One HMC function can manage multiple IBM Z CPCs. When tasks are performed at the HMC function, the commands are routed to the Primary SE function of the IBM z17. Then, the SE function issues those commands to the targeted CPC.

For more information about feature codes for the supported physical HMCs, see 10.1 "HMC and SE introduction", in the *IBM z17 Technical Guide*, SG24-8579.

HMC users

Password rules for all default users changed. In 2018, the state of California passed the "password law" that bans the use of default passwords in connected devices (see California SB-327). It requires that any default passwords that are included with connected devices are changed on installation or are unique per device.

These rules are implemented by IBM on all IBM z17 devices worldwide. Therefore, a first logon to the default user forces a password logon change. The client is responsible for maintaining the passwords.

Note: Default users that are predefined on the IBM z17 machines are ACSADMIN and SERVICE. Former users such as ADVANCED, OPERATOR, SYSPROG, STORAGEADMIN and other individual users can be re-created by using the HMC ACSADMIN user ID.

However, default user roles ADVANCED, OPERATOR, STORAGEADMIN, and SYSPROG continue to be included, so you can create user IDs from them.

Because the password for the SERVICE default user must be changed, every client must establish a plan for the authorized access to the password for this user for the following reasons:

- The IBM Systems Service Representative (IBM SSR) might be a different person for various visits.
- IBM SSRs might show up at any time (including the middle of the night) for planned (maintenance actions) or unplanned (repair actions) activities that involve an IBM Z server.

The client should be ready to provide the user ID and the password to the IBM SSR on arrival on the site.

- The client should maintain a list of unique Service IDs and passwords.
- The client must establish a process for the hand over of user IDs and passwords to the IBM SSR to avoid service delays.

The configure data replication task

Data replication between the various HMCs is a simple and effective way to keep all the parameters and settings in sync.

As data is replicated from one HMC function to another, an internal level indicator for the data that is replicated is incremented each time that the data is altered on the data source. Each HMC function tracks the level indicator for each type of data and does not accept data from a data source when the level indicator is not greater than what is on the receiving HMC function.

The HMC data replication service setup is a wizard-guided process on the HMC function. When logged on as an ACSADMIN, click **Task Index** on the left of the window, and then click **Configure Data Replication**. Figure 2-5 on page 23 and Figure 2-6 on page 24 show the entry level for setting up the data replication task.

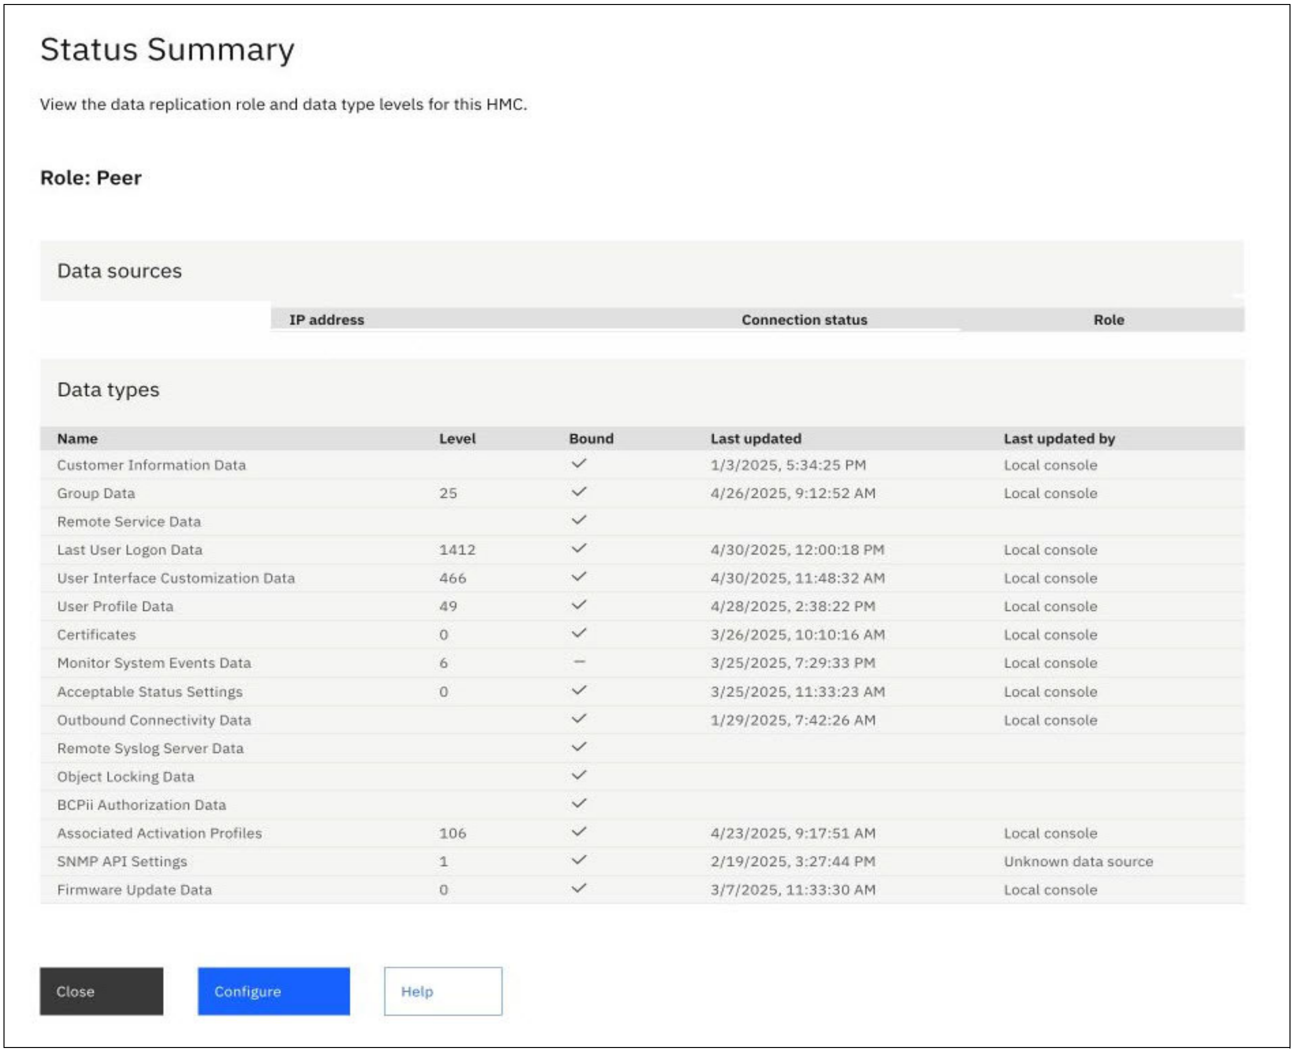

Figure 2-5 Entry window for configuring an HMC data replication task

Follow the windows and complete the setup of the data replication task.See Figure 2-6 on page 24.

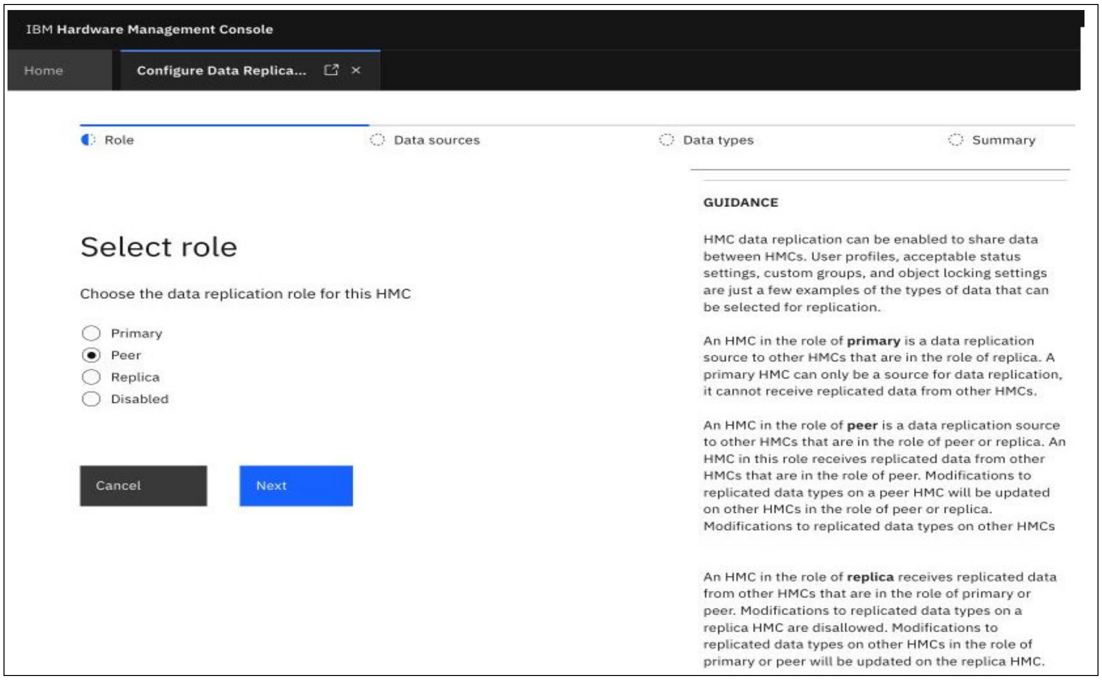

Figure 2-6 HMC role selection for replication

2.4.7 Support Element settings

The HMAs that are supplied with the IBM z17 are two appliances that are supported by the 1U HMA servers. Both SE server units are installed at the top of the A frame. One runs the primary SE function, and the other runs the alternate SE function.

Generally, the SE settings are considered part of the physical installation of the IBM z17, so they are not presented in this book. For a new IBM z17 system, a new range of TCP/IP addresses must be provided by the customer to the IBM SSR who performs the physical installation. As an extra measure of security, provisioning of a separate LAN segment for the management functions is preferred. During an upgrade from a previous generation IBM Z platform to an IBM z17, the current settings on the SEs should be backed up for migration purposes.

In addition to the standard SE configuration, there might be other parameters that should be backed up, such as the API settings. These parameters can be accessed through the Customize API Settings task on the SE function.

Any default user IDs that are part of a previous HMC level can be carried forward to the new HMC levels as part of a MES upgrade or by selecting “User profile data for the Save / Restore customizable console data” or “Configuration Data Replication tasks” using the ACSADMIN role on the HMC function.

Figure 2-7 shows the main selection window on the SE.

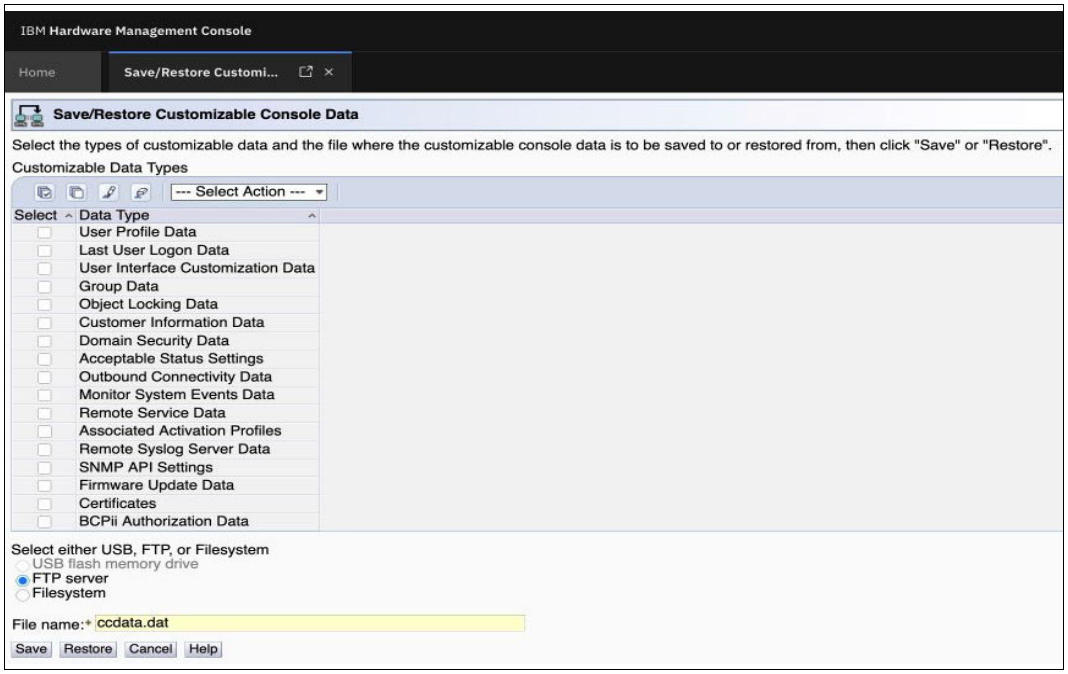

Figure 2-7 Save/Restore Customizable Console Data

In Figure 2-7, the option to save data to an FTP server is selected. At the time of writing, the USB Interfaces on the 1U SE servers were available to the client. However, it is a best practice to provide access to an FTP server for all the clients that are attached to the HMC LAN.

2.4.8 Network Time Protocol

The Network Time Protocol (NTP) is a networking protocol for clock synchronization between computer systems over packet-switched, variable-latency data networks. In operation since 1985, NTP is one of the oldest Internet protocols in current use. NTP was designed by David L. Mills of the University of Delaware.

NTP is intended to synchronize participating computers to within a few milliseconds of Coordinated Universal Time (UTC). It uses the intersection algorithm, a modified version of Marzullo's algorithm[4], to select accurate time servers and is designed to mitigate the effects of variable network latency.

2.4.9 Precision Time Protocol

The Precision Time Protocol (PTP) standard enables accurate and precise synchronization of the real-time clocks of devices in networked distributed systems. The protocol is applicable to systems where devices communicate through networks, including Ethernet. The standard allows multicast communication, unicast communication, or both.

PTP enables heterogeneous systems that include clocks of various inherent precision, resolution, and stability to synchronize to a grandmaster clock. The protocol supports

[4] Marzullo's algorithm, invented by Keith Marzullo in 1984, is an agreement algorithm used to select sources for estimating accurate time from a number of noisy time sources. NTP uses a modified form of this called the Intersection algorithm, which returns a larger interval for further statistical sampling.

synchronization in the submicrosecond range with minimal network bandwidth and local clock computing resources. The protocol enhances support for synchronization to better than1 nanosecond.

2.4.10 Server Time Protocol

Server Time Protocol (STP) is a time synchronization architecture that is designed to provide the capability for multiple servers to maintain time synchronization with each other and to form a Coordinated Timing Network (CTN). STP is a message-based protocol that allows timekeeping information to be sent between servers and Coupling Facilities over:

- ICA SR (Integrated Coupling Adapter Short Reach) Channel Type CS5 links IBM z15™ and IBM z16™ only.
- ICA SR1.1 (Integrated Coupling Adapter Short Reach) links IBM z15, IBM z16 and IBM z17.
- ICA SR2.0 (Integrated Coupling Adapter Short Reach) links IBM z15, IBM z16 and IBM z17.
- Coupling Express2 Long Reach (CE LR), channel type CL5 links IBM z15, IBM z16. It also links IBM z17 equipped with Coupling Express3 Long Reach 10G.
- Coupling Express3 Long Reach 25G, channel type CL6 links to other Coupling Express3 Long Reach 25G only.

STP operates with the TOD-clock steering facility to provide a new timing mode, timing states, external interrupts, and machine check conditions.

2.4.11 Pulse per second

Since the introduction of IBM Server Time Protocol (STP) for IBM Z in 2005/2006, the Network Time Protocol (NTP) was supported as an external time reference for STP. When the accuracy capability of NTP alone was not good enough to comply with applicable industry regulations, IBM Z clients could supplement NTP with pulse per second (PPS) capability via coaxial cable connection from the same time server providing the NTP connection.

Additional information can be found in this IBM White Paper.

2.4.12 External Time Sources (ETS) and STP planning considerations

The external time sources (ETS) NTP, PTP and PPS that have been described in the previous sections are used for the Central Processor Complex (CPC) to keep the time on the CPC as accurate as possible.

Up to the predecessor machine IBM z16, ETS was restricted to two time sources of either NTP or PTP, and the CPC would use the primary of the two sources as long as it was functional and would only switch to the backup when the primary ETS became unavailable.

Each CPC Drawer of an IBM z17 includes two oscillator cards for time services, and each oscillator card has one port for PPS and one port for ETS (NTP and/or PTP) as shown in Figure 2-8.

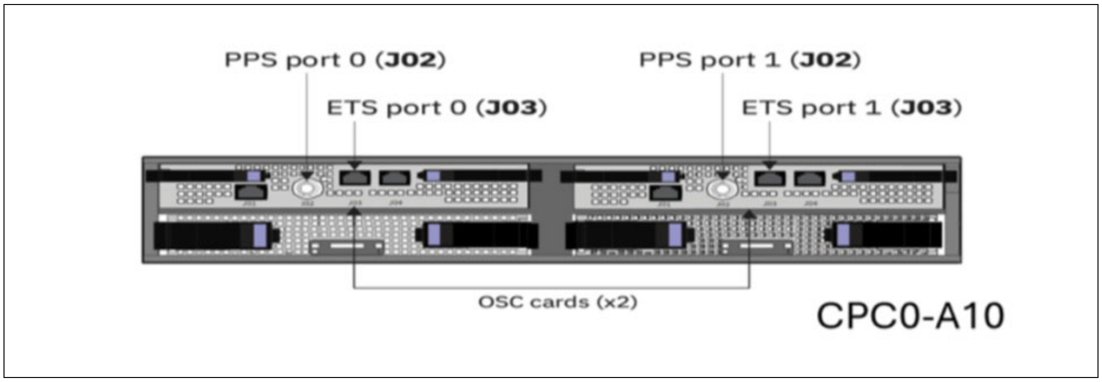

Figure 2-8 ETS ports of a CPC Drawer

Beginning with IBM z17, you are allowed to connect up to three NTP sources and in addition up to two PTP sources to the CEC. But still only two time source cables are allowed to be connected to an IBM z17 for NTP and PTP though, even if there are more CPC Drawers and by that more ETS ports available in the CPC. Therefore you need switches to connect more than two external time sources of the types NTP or PTP to the CPC. PPS is connected through its own coax PPS port.

Note: For accuracy reasons, IBM recommends the concurrent use of either NTP combined with PPS, or NTP in combination with PTP for optimal results. If possible omit PPS in favor of PTP for ease of use.

The CPC will take all connected time sources into account to acquire the most accurate time possible. The advantage of NTP is, that this protocol can be secured using Network Time Security (NTS), while the advantage of PTP is the higher accuracy.

By combining both types and multiple sources, you achieve the highest reliability and accuracy of external time source. Figure 2-9 on page 28 shows an example of how multiple time sources can be connected to the CPC. If the CPC only has one CPC Drawer, all ETS cabling would go to this drawer instead of being spread over two drawers.

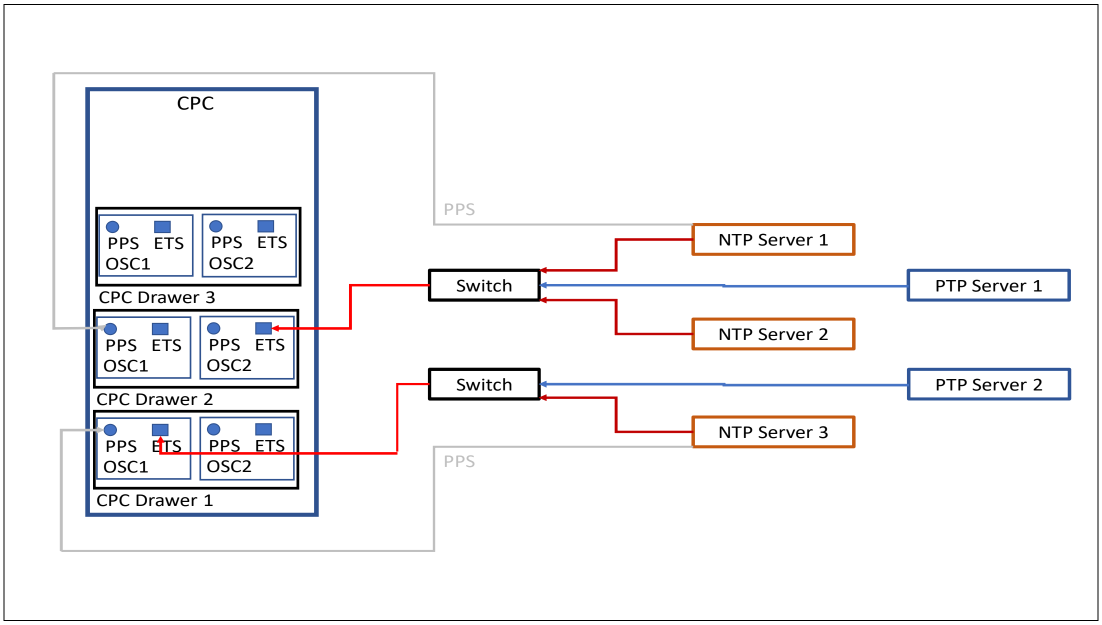

Figure 2-9 Connecting multiple NTP and PTP sources to a CPC

By combining both types of protocols and multiple sources of each, you achieve the highest reliability and accuracy of external time source.

IBM Z servers implement a Coordinated Time Network (CTN) using Coupling Links to synchronize time between multiple IBM Z servers. One exploiter of the CTN is Sysplex Time Protocol (STP). Setting up CTN and STP is described in more detail in Chapter 8, “Preparing for IBM Parallel Sysplex and Server Time Protocol” on page 153

.

CTN connectivity for an IBM z17 system: Timing data is exchanged between IBM Z servers by using coupling links. An IBM z17 system can connect to the IBM z15 server family and the IBM z16 server family as well as additional IBM z17 servers with compatible coupling links (N-2 generations coupling links). Make sure that you have the appropriate coupling and timing connectivity in your IBM z17 system before assigning a role in the Coordinated Timing Network (CTN) (Preferred Time Server (PTS), Backup Time Server (BTS), and Arbiter).

Note: Since IBM z15, support for STP stratum level 4 is provided. This feature helps avoid the extra complexity and expense of system reconfiguration. All systems that might become exposed to this situation should have this change installed. Stratum level 4 should be used only during a migration, and for a short period.

Although STP stratum level 4 is supported, it should not be used for permanent configurations. Stratum 4 should be used for transitional configurations during CTN maintenance.

For more information, see Chapter 8, “Preparing for IBM Parallel Sysplex and Server Time Protocol” on page 153.

2.5 IODF configuration

This section describes I/O configuration considerations in the IODF.

2.5.1 Channel subsystems

The IBM Z platform manages I/O resources (LPARs, channel paths, control units (CUs), and I/O devices) by housing them in multiple Logical Channel Subsystems (LCSSs).

A *spanned channel path* is a channel that can be used by partitions in more than one LCSS. Use the same CHPID value across all LCSSs sharing a spanned channel. However, LCSSs that do not share a spanned channel can use that CHPID for other channels.

For more information, see *z/OS Hardware Configuration Definition Planning*, GA32-0907.

> **Note:** In this book the designation Channel Subsystem (CSS) and Logical Channel Subsystem (LCSS) are used synonymously.

Your planning should consider multiple LCSSs so that you can logically partition your physical channel resources to accommodate large-scale enterprise workload connectivity and high-bandwidth demands. IBM z17 ME1 supports six LCSSs and has four Subchannel Sets (SSs) in each LCSS with up to 256 channels, for a total of 1536 channels.

Also, LCSSs provide for multiple SSs for expanding the number of I/O devices that are managed in each CSS. With IBM z17 ME1, up to four SSs supporting 65280 devices in SS 0 and 65535 devices for each other SS. The DASD base addresses are defined to SS0. IBM reserves 256 subchannels on set 0 and the alias addresses are either to SS 1, 2, and 3.

Not all device types are eligible for nonzero SSs. SS0 can be used for any type of device. More SSs (for example, SS1) can be used only for certain classes of devices, such as parallel access volume alias devices.

For more information, see *IBM z17 (9175) Technical Guide,* SG-24-8579. Use multiple SSs to move devices of eligible device types to more SSs, and then define more physical devices to SS0.

2.5.2 Logical partitions

With the PR/SM feature, a single IBM Z can run multiple OSs and Coupling Facilities (CFs) in logical partition (LPAR) mode. Each OS and each CF has its own LPAR, which contains a separate set of system resources that includes these items:

- A portion of storage (memory).
- One or more central and specialty processors. The processors can be dedicated or shared.

Profile data can be exported on the older IBM Z platform and imported on the IBM z17 system. If the LPAR data is imported from an older IBM Z platform, consider the LPAR sizing before the LPAR migration to the IBM z17 system. For more information, see *Support Element Operations Guide* (link requires an IBM Resource Link valid user ID to access).

For more information about how to define LPARs in IODF, see Chapter 3, “Preparing for a new IBM z17 system” on page 39.

Planning considerations for Virtual Flash Memory

VFM (Feature Code 0566) is available in 512 GB increments of memory. IBM z17 ME1 can have up to 12 VFM features. While planning your memory, you must consider your VFM requirements.

With the introduction of VFM, there are no changes to the existing OS interface for handling the storage-class memory (SCM). The allocation of VFM storage is done during LPAR activation because the LPAR hypervisor manages the partition memory.

Both the initial and maximum amounts of VFM are specified in the LPAR image profile. VFM can be added to or deleted from OSs by using existing SCM commands after the LPAR is activated. VFM allocation and definition for all partitions can be displayed on the Storage Information window on the HMC and by using SCM commands in z/OS.

VFM allocation: The VFM values for Initial and Maximum allocations cannot be dynamically changed. One or more partitions must be activated (or reactivated) for VFM allocation changes to take effect.

As a best practice, assign the maximum amount that is installable for all LPARs that are candidates for using VFM, and set Initial allocation to 0 (zero) for the LPARs that do not require immediate activation of VFM. By doing so, you ensure that you can later use any available VFM when required.

At partition activation time, over-commitment of VFM storage is supported. This setting enables more storage to be added to partitions that is subject to the amount that is not assigned to other partitions.

For more information, see 10.2.3, “Configuring VFM” on page 224.

If the total amount of VFM that is allocated to all active partitions is equal to the LICC value but the sum of active partition maximums is larger than the installed amount, the client might be able to concurrently add VFM and increase allocations without reactivating partitions. This feature is illustrated in the following sections.

Non-disruptive migration

Here is an example of a non-disruptive migration:

- An IBM z17 has three VFM features that are installed (512 GB each), with a LICC value of 1.5 TB.
- LPAR A has 1.0 TB that is assigned, with a maximum value of 1.5 TB.
- LPAR B has 512 GB that is assigned, with a maximum value of 1.0 TB.
- Assign 1 TB to LPAR B, but this change is not possible within the constraints of the installed VFM.
- You can purchase and install another 512 GB VFM feature and install it concurrently. Now, up to 512 GB can be added concurrently to LPAR B without reactivating the LPAR.

Figure 2-10 shows the non-disruptive migration example.

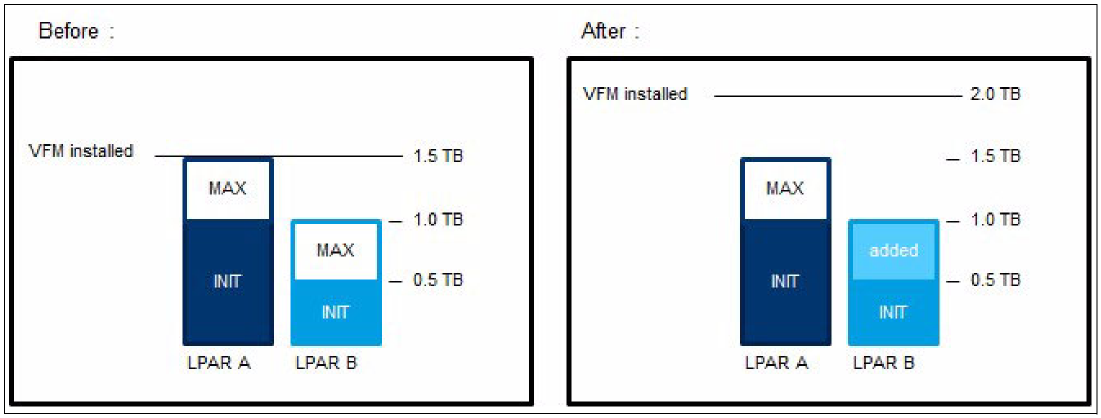

Figure 2-10 Non-disruptive VFM migration example

Disruptive migration

Here is an example of a disruptive migration:

- An IBM z17 has two VFM features that are installed (512 GB per feature), with a LICC value of 1.0 TB.
- LPAR A has 512 GB that is assigned, with a maximum value of 1.0 TB.
- LPAR B has 256 GB that is assigned, with a maximum value of 1.0 TB.
- Change LPAR A so that it can have up to 1.5 TB. This change falls outside the range of the maximum installed VFM.
- You can purchase and install two more 512 GB VFM features concurrently (assuming that the memory is ordered and available). Now, you must reactivate LPAR A with the new maximum VFM value of at least 1.5 TB and less than or equal to 2.0 TB.

> **Note:** Plan-Ahead Memory is not available on the IBM z17.

Figure 2-11 shows the disruptive VFM migration example.

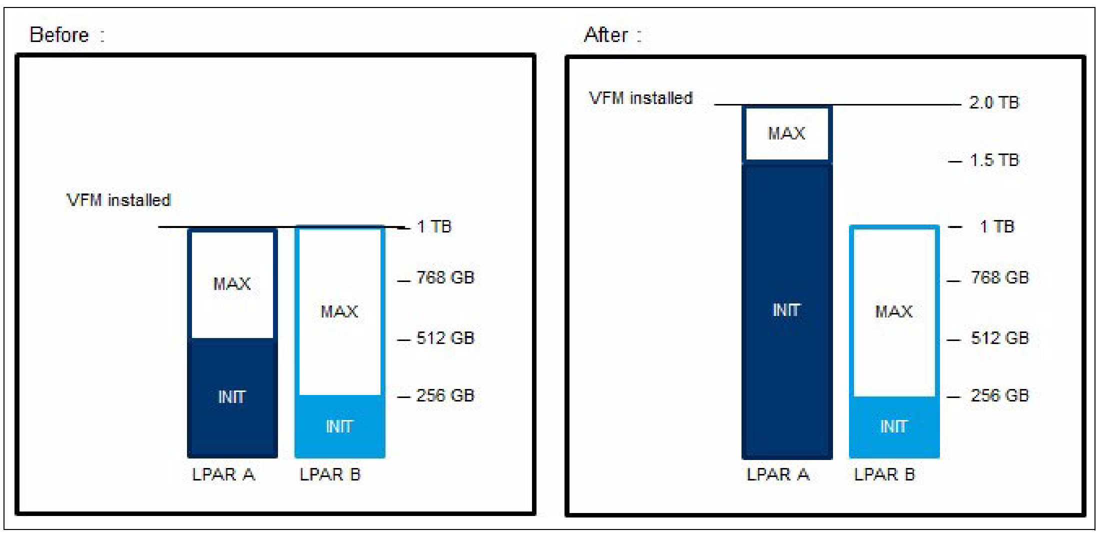

Figure 2-11 Disruptive VFM migration example

For more information about how to configure VFM, see 10.2, "Virtual Flash Memory" on page 223.

2.5.3 Storage connectivity

The FICON Express32-4P and FICON Express16SA (upgrade only) provide connectivity to storage devices by using FICON or FCP protocols. FICON Express32-4P supports negotiation to 8, 16, or 32 gigabits per second (Gbps) link data rates. 8 Gbps transfer rates can be achieved through a FICON Director with 8 or 16 Gbps Optics. FICON Express 16SA supports 8 and 16 Gbps data link rates.

The FICON Express features support IBM High Performance FICON for IBM Z (zHPF). zHPF is an extension to the FICON architecture that provides performance improvement for single-track and multi-track operations.

On a new IBM z17 system, only the FICON Express32-4P features can be ordered. FICON Express16SA can be carried forward when upgrading from an older IBM Z.

> **Note:** With the FICON Express32-4P and FICON Express16SA, *all* ports must be configured either as channel type FC or FCP. A mixed configuration is *not* allowed.

For more information about how to configure the FICON Express32-4P feature, see Chapter 12, “Adding storage devices” on page 243.

IBM zHyperlink Express

zHyperLink Express is a short-distance IBM Z I/O channel that is designed for up to 10x lower latency than zHPF. This feature is in the PCIe+ I/O drawer, and it is a 2-port card that is used for short-distance direct connectivity between an IBM z17 system and an IBM DS8000® system. A customer-supplied 24x MTP-MTP cable is required for each port of the zHyperLink Express feature.

zHyperLink Express is designed for distances up to 150 m and supports a link data rate of 8 GBps. A zHyperlink port is fully sharable between all partitions because 127 virtual functions (VFs) / PCIe function IDs (PFIDs) per link are supported.

IBM zHyperLink reduces the latency of direct access storage device (DASD) I/Os by interconnecting the IBM z17 system directly to the I/O bay of a DS8880 or later storage system. This feature improves the application response time without application changes. zHyperLink is fast enough to run I/Os synchronously so that the CPU can wait for the data, which results in the following advantages:

- No undispatch of the running task
- No CPU queuing delays to resume it
- No host CPU cache disruption
- Small I/O service time

The zHyperLink Express adapter takes one slot on an IBM z17 PCIe+ I/O drawer, and each adapter has a single PCHID with two ports. Up to 16 zHyperLink Express adapters can be installed in one IBM z17 system, which means that you can have up to 32 ports.

FICON connectivity to each storage system is still required for these purposes:

- For initialization of the zHyperLink connection
- For I/Os that are not eligible for zHyperLink
- For fallback when a zHyperLink request fails (for example, a cache miss or busy condition)

zHyperLink reads are exploited by Db2 for synchronous reads and VSAM for reads (except NSR sequential). zHyperLink writes are exploited by Db2 for the active log. zHyperLink is also supported by IBM MQ. It can improve active log throughput and reduce IBM MQ transaction time by up to 3.5 times.

For more information about the zHyperLink feature, see *IBM z17 (9175) Technical Guide*, SG24-8579, *IBM z16 (3931) Technical Guide,* SG24-8951, *IBM z16 A02 and IBM z16 AGZ Technical Guide,* SG24-8952, *IBM Z Connectivity Handbook,* SG24-5444, and *Getting Started with IBM zHyperLink for z/OS,* REDP-5493.

For more information about defining zHyperLink Express, see 15.2.4, “Defining a zHyperLink PCIe function” on page 345.

For more information about zHyperLink Express management, see “Managing zHyperLink Express” on page 348.

2.5.4 Network connectivity

This section provides planning considerations for deploying the following network-related features:

- Network Express Adapter
- OSA-Express
- Shared Memory Communication (SMC)
  - SMC-R
  - SMC-D
- HiperSockets

Network Express Adapter

Starting with IBM z17, the new Network Express Adapter is supported. It combines the functionalities:

- OSH Channel type: Supports all legacy functions available with OSD over OSA-Express but uses Enhanced Queued Direct I/O (EQDIO) architecture.
- NETH Function type: Supports SMC-R communication.

> **Note:** IBM z17 no longer supports RoCE Express adapters

The Network Express Adapter has two ports. Each of them is represented by its own PCHID. Each of the two PCHIDs is capable to be defined for OSH and NETH at the same time. By this, a consolidation of network features can be achieved.

The operation mode of the Network Express Adapter for OSA traffic is OSH, which uses enhanced QDIO (eQDIO). It is the equivalent to OSD (using QDIO) on the OSA Express card.

> **Note:** eQDIO is not supported for Linux on Z running natively in an LPAR.
>
> If you run Linux on Z under z/VM, you can define OSH to a VSWITCH in z/VM and connect the Linux to this VSWITCH.

The Network Express Adapter is offered in four variations:

- Network Express 10G LR (FC0525) and SR (FC0524)
- Network Express 25G LR (FC0527) and SR (FC0526)

Definition and configuration of OSH and NETH do not require any actions in OSA Advanced Facilities on the HMC. The details of defining OSH and is described in 13.4, “Network Express for channel type OSH” on page 287 and for NETH in 15.2.3, “Defining a NETH PCIe function” on page 339.

OSA-Express

The OSA-Express features are installed into an IBM z17 PCIe+ I/O drawer. The features are available as different types and support different networking protocols.

> **Note:** The operation mode OSE is no longer supported on IBM z17. Additional planning is needed to migrate from OSE to OSD.

Based on the intended usage, the operating modes must be defined with a channel type and device address. For configuration details of operation mode OSD, see Chapter 6, “Configuring network features” on page 129. For configuration details of operation mode

OSC, see Chapter 7, “Defining console communication” on page 137. Additional details for both operation modes can be found in the *OSA-Express Implementation Guide,*SG24-5948.

Starting with Driver Level 22 (HMC V2.13.0), HMC is enhanced to leverage the OSA Advanced Facilities. For all supported OSA-Express features, OSA Advanced Facilities on the HMC is required. The OSA Advanced Facilities is used primarily for these purposes:

- Manage all OSA-Express ports.
- Configure all OSA-Express OSA-ICC ports.
- Configure local Media Access Control (MAC) addresses.
- Display OSA address table to verify registered Internet Protocol Version 4 (IPv4) addresses (in use) for QDIO ports.
- Display registered IPv4 or IPv6 Virtual MAC addresses and virtual local area network IDs that are associated with all OSA-Express features that are configured as QDIO Layer 2.
- Provide status information about an OSA-Express port.

For more information about the use of OSA Advanced Facilities on the HMC, see Chapter 6, “Configuring network features” on page 129.

The OSA Express feature is available in following variations:

- OSA Express 7S 1000BaseT (FC 0446) [5,6]
- OSA Express 7S 1GbE LX (FC0442) and SX (FC0443)
- OSA Express 7S 10G LR (FC0444) and SR (FC0445)
- OSA Express 7S 1.2 1GbE LX (FC0454) and SX (FC0455)
- OSA Express 7S 1.2 10G LR (FC0456) and SR (FC0457)
- OSA Express 7S 1.2 25G LR (FC0460) and SR (FC0459)

Shared Memory Communication (SMC)

The SMC protocol is designed to help reduce CPU consumption, reduce latency and enlarge throughput for TCPIP traffic between LPARs on IBM Z platforms. With SMC, after a TCPIP connection has been established over standard communication channels (OSD, OSH or Hipersockets), both endpoints of the communication can negotiate on using SMC for the ongoing transmissions of this connection, if a suitable SMC infrastructure is available and associated to the initial communication channel. SMC is transparent to the applications.

Using the TCPIP protocol, each transmission has to transition through all TCP/IP layers on both endpoints. This effort is spared when using SMC which reduces TCP/IP CPU and latency usage.

SMC is available in two versions:

- SMC version 1 connectivity is limited to hosts that are directly attached to a common single IP subnet. It is purely Layer 2 communication. In addition the same PNETID has to be defined for the SMC function as well as to the associated OSA or a Network Express Adapter.
- SMC version 2 reuses the existing SMCv1 (CLC and LLC) messages, but it does not require a PNETID for SMC-D to be defined for the SMC function or the associated OSA. In addition, SMCv2 traffic is not limited to a common single IP subnet, but can be routed.

For more information about the SMC message (wire) flows, see RFC 7609.

[5] only available as carry forward from IBM z15
[6] only for operation mode OSD, not for OSC

To exploit SMC, additional configuration steps have to be taken at the TCPIP stack. See *z/OS Communications Server: IP Configuration Guide*, SC27-3650 for more details.

Before implementing SMC-R or SMC-D, check your environment for the following items:

- Run SMCAT to evaluate its applicability and potential value. For more information about SMCAT, see *IBM z/OS SMC Applicability Test (SMCAT)* and 2.3.5, “Shared Memory Communications Applicability Tool” on page 15.
- Review and adjust the available real memory and fixed memory usage limits (z/OS and CS) as needed. SMC requires fixed memory. Review the limits and provision more real memory for z/OS.
- Review the IP topology, VLAN usage considerations, and IPsec (not supported).
- Review changes to messages, monitoring information, and diagnostic tools. There are numerous updates to these items:
  - Messages (IBM Virtual Telecommunications Access Method (IBM VTAM®) and TCP stack)
  - The `netstat` command (status, monitoring, and display information)
  - CS diagnostic tools (VIT, Packet trace, CTRACE, and IPCS formatted memory dumps)

SMC-R

SMC using RDMA over Converged Ethernet (RoCE) is called SMC-R. It allows SMC communication between LPARs of distant CPCs. It is provided by the NETH function available with the Network Express Adapter. The NETH function has to be associated with either an OSH channel on the same PCHID as the NETH function, or with an OSD channel on an OSA Express card.

Deployment of the NETH features is supported in either point-to-point or switched configurations. When you plan to deploy NETH features in a switched configuration, the switches must support the following requirements:

- Global Pause function frame (as described in the IEEE 802.3x standard) should be enabled.
- Priority Flow Control (PFC) should be disabled.

The definition of NETH is described in more detail in 15.2.3, “Defining a NETH PCIe function” on page 339.

SMC-D

SMC using the Internal Shared Memory (ISM) is called SMC-D. It allows SMC communications between LPARs within the same CPC.

SMC-D uses ISM virtual PCIe (vPCIe) adapters to provide Direct Memory Access (DMA) communications between LPARs inside the same IBM Z server.

SMC-D is a protocol that allows TCP socket applications to transparently use ISM. ISM is a virtual channel like Internal Queued Direct (IQD) for HiperSockets. A virtual adapter is created in each z/OS LPAR. By using the SMC protocol, the memory is logically shared. The virtual network is provided by firmware.

The definition of ISM is described in more detail in 15.2.2, “Defining an ISM PCIe function” on page 334.

HiperSockets

The HiperSockets function is based on the OSA-Express QDIO protocol, so HiperSockets is called *internal queued direct input/output* (IQDIO). LIC emulates the link control layer of an OSA-Express QDIO interface and uses no physical cabling or external networking connections. Data access is performed at memory speeds, bypassing external network delays and providing users with high-speed logical LANs with minimal system and network impact.

HiperSockets can be defined as MIF-shared in a CSS and as spanned channels across multiple CSSs. A HiperSockets CHPID can be seen as an *internal LAN* to the server. The level of sharing is determined by the LPARs that you want to grant access to that LAN.

HiperSockets are supported by the following OSs:

- All in-service z/OS releases
- All in-service z/VM releases
- 21st Century Software VSE^n 6.3.1 and higher
- Linux on IBM Z:

The definition of Hipersockets is defined in more detail in 13.5, “IQD CHPIDs for HiperSockets” on page 292.

For more information about the technical details of HiperSockets function, see: *IBM Z Connectivity Handbook*, SG24-5444.

2.5.5 Coupling and timing links

Support for Parallel Sysplex includes the Coupling Facility Control Code (CFCC) that runs in a LPAR, and coupling links. Coupling connectivity in support of Parallel Sysplex environments is provided on the IBM z17 by the following features:

- Coupling Express3 LR 10G (Feature Code 0498) has 2-port coupling link connectivity for a distance up to 10 km (6.2 miles) and is compatible with the CE LR and the CE2 LR links from N-1 or N-2 generation IBM server.
- Coupling Express3 LR 25G (Feature Code 0499) has 2-port coupling link connectivity for a distance up to 10 km (6.2 miles) and is compatible with the CE3 LR only.
- Integrated Coupling Adapter Short Reach (ICA SR 2.0) (Feature Code 0216) which is compatible with ICA SR and ICA SR1.1.
- Internal Coupling (IC) channels operate at memory speed.

The number of physical coupling links and logical coupling CHPIDs that is supported by each IBM z17 can be found in the*IBM Z Connectivity Handbook*, SG24-5444.

All coupling link types, except IC links, can be used to carry STP or PTP messages.

Note: The CE3-LR is a 2-port card that occupies one PCIe+ I/O drawer slot. Therefore, an IBM z17 that is configured as a Stand-Alone Coupling Facility (SACF) must have at least one PCIe+ I/O drawer for LR coupling.

Planning consideration

The relationship between one or more CF link connections among CPCs must be configured in HCD to enable the exchange of CF link signals. HCD generates the CU and device definitions automatically if the CPCs are known within the same IODF file and the adapter ID (AID) or PCHIDs are not reserved by other definitions.

> **Coupling connectivity for IBM z17:** IBM z17 supports coupling connectivity back to IBM z15 T01 and IBM z15 T02. Coupling is supported only through ICA SR, ICA SR 1.1, CE LR and CE2 LR links.
>
> N-3 connectivity for coupling links is not supported.
>
> To manage an IBM z17 system in a CTN, HMC level 2.17.0 or later must be used. The IBM z17 SE no longer supports STP menus.

As described in this section, depending on the type of the CF link hardware, CF links operate up to a set distance. Physical placement of the CPCs or CFs must be considered to avoid exceeding the maximum distance that is supported by the CF link. For the CE3-LR links, Dense Wavelength Division Multiplexing (DWDM) technology can be used to extend the maximum length of the CF links.

For a list of qualified devices, log in to the IBM Resource Link website with a registered Resource Link ID.

STP or PTP signals can be exchanged between two CPCs without any CF LPARs that are involved. If physical coupling links are established between two CPCs, HCD enables the configuration of STP links (timing-only links).

For more information, see *z/OS HCD User's Guide*, SC34-2669, and Chapter 8, "Preparing for IBM Parallel Sysplex and Server Time Protocol" on page 153.

2.5.6 Planning considerations for hardware data compression

This section provides planning considerations for enabling hardware data compression on an IBM z17 system.

The IBM z17 Tellum II chip design incorporates a compression unit, which is the IBM Integrated Accelerator for z Enterprise Data Compression (zEDC), and a Compression Coprocessor (CMPSC) that is implemented on every Processing Unit (PU) core. Improved data compression operations are achieved by having one dedicated compression coprocessor per PU, and new hardware instructions.
The compression coprocessor and the IBM Integrated Accelerator for zEDC use algorithms for data compression that enables the reduction in the size of data to save storage space or increase the data transfer throughput.

Here is a short summary about planning considerations for hardware data compression:

1. Planning the installation:
   - Update the IFAPRDxx PARMLIB member in the z/OS operating system.
   - Plan for IPLs before activating the software feature for the first time.
2. z/OS: Verifying the prerequisites: Look up the*IBM.Function.zEDC* FIXCAT for proper PTFs.
3. z/OS: Enabling the priced software feature provides native compression support for the use on z/OS libraries eligible for compression by zEDC. The zlib Java library is provided by IBM "as is".

3

Preparing for a new IBM z17 system

This chapter describes two scenarios when preparing for an IBM z17 installation:

- Upgrading an existing IBM z15 T01 or an IBM z16 A01 to an IBM z17 ME1 while maintaining your existing serial number. An upgrade includes the frames, CPC drawers (central processor complex and PCIe+ I/O drawers), and new build or carry-forward I/O features.
- Installing a new IBM z17 ME1 into an existing environment.

Because many environments exist, the results that are achieved in your environment might differ from the ones that are described here.

Naming: The IBM z17 system that is targeted by this publication consist of an IBM z17 ME1. Throughout this chapter, we might refer to this machine as IBM z17.

This chapter includes the following topics:

- Supported hardware features
- Saving and restoring Open Systems Adapter-Express configuration data
- Upgrading an IBM z16 to an IBM z17 while maintaining the serial number
- To define the I/O configuration for your system, go to Chapter 5, “Building the production input/output definition file and setting up the central processor complex” on page 89.Installing a new IBM z17 system into an existing IBM Z environment

3.1 Supported hardware features

This section lists the channel path ID (CHPID) types and hardware features for the IBM z17 ME1 (9175).

The CHPID types:

- OSH (for the Network Express Adapter)
- FC / FCP FICON (Express32-4P Port card are introduced for IBM z17)
- CL6

Here are the hardware features that are new or available for order:

- Feature Code 0216 - Integrated Coupling Adapter SR2.0
- Feature Code 0351 - zHyperLink Express2.0
- Feature Code 0387 - FICON Express32G-4P LX
- Feature Code 0388 - FICON Express32G-4P SX
- Feature Code 0454 - OSA-Express7S 1.2 GbE LX
- Feature Code 0455 - OSA-Express7S 1.2 GbE SX
- Feature Code 0456 - OSA-Express7S 1.2 10GbE LR
- Feature Code 0457 - OSA-Express7S 1.2 10GbE SR
- Feature Code 0459 - OSA-Express7S 1.2 25GbE SR
- Feature Code 0460 - OSA-Express7S 1.2 25GbE LR
- Feature Code 0498 - Coupling Express3 LR 10Gb
- Feature Code 0499 - Coupling Express3 LR 25Gb
- Feature Code 0524 - Network Express SR 10G
- Feature Code 0525 - Network Express LR 10G
- Feature Code 0526 - Network Express SR 25G
- Feature Code 0527 - Network Express LR 25G
- Feature Code 0566 - IBM Virtual Flash Memory (VFM)
- Feature Code 0909 - Crypto Express8S (1 port)
- Feature Code 0908 - Crypto Express8S (2 ports)

Here are the CHPID types that can be migrated (carry forward):

- Fibre Channel (FC) and Fibre Channel Protocol (FCP)
- OSC and OSD
- CS5, CL5, and ICP
- Internal Queued Direct (IQD)

Here are the hardware features that can be migrated (carry forward).

> **Note:** When migrating from an IBM z15 to an IBM z17, all defined CHPIDs with type OSM have to be deleted without any replacement.
>
> CHPID type OSE is no longer supported on IBM z17.

> **Note:** OSA Express7S 1000BaseT is not supported for CHPID type OSC on an IBM z17. OSC requires an OSA Express7S 1.2 GbE SX/LX Adapter card.
>
> OSA Epress7S 1000BaseT can be carried forward for CHPID type OSD only from an IBM z15.

- Feature Code 0446 - OSA-Express7S 1000Base-T Ethernet (IBM z15 only)

- Feature Code 0442 - OSA Express 7S GbE LX (IBM z15 only)
- Feature Code 0443 - OSA Express 7S GbE SX (IBM z15 only)
- Feature Code 0444 - OSA Express 7S 10GbE LR (IBM z15 only)
- Feature Code 0445 - OSA Express 7S 10GbE SR (IBM z15 only)
- Feature Code 0454 - OSA-Express7S 1.2 GbE LX
- Feature Code 0455 - OSA-Express7S 1.2 GbE SX
- Feature Code 0456 - OSA-Express7S 1.2 10GbE LR
- Feature Code 0457 - OSA-Express7S 1.2 10GbE SR
- Feature Code 0459 - OSA-Express7S 1.2 25GbE SR
- Feature Code 0460 - OSA-Express7S 1.2 25GbE LR
- Feature Code 0898 - Crypto Express7S (1 port)
- Feature Code 0899 - Crypto Express7S (2 ports)
- Feature Code 0436 - FICON Express16SA LX
- Feature Code 0437 - FICON Express16SA SX
- Feature Code 0461 - FICON Express32S LX
- Feature Code 0462 - FICON Express32S SX

For more information about the supported I/O features, see *IBM Z Connectivity Handbook*, SG24-5444.

3.2 Saving and restoring Open Systems Adapter-Express configuration data

The two processes for OSA cards that you might need to use when upgrading or replacing your processor are described here.

Exporting and importing OSA-ICC configuration data with OSA Advanced Facilities

If you are unfamiliar with the exporting and importing process for OSA Integrated Console Controller (OSA-ICC) Server and Session configuration data, see 7.3, “Configuring a new OSA-ICC configuration by using OSA Advanced Facilities” on page 138.

Using OSA Advanced Facilities to set OSA parameters

For more information about the process of changing the Media Access Control (MAC) addresses, see 6.1, “Configuring OSA for OSA-Express and Network Express Adapter” on page 129.

3.3 Upgrading an IBM z16 to an IBM z17 while maintaining the serial number

This section describes the steps to upgrade an existing IBM z16 that is defined in your input/output definition file (IODF) to an IBM z17 and keeping the system serial number.[1]

[1] For experienced Hardware Configuration Definition (HCD) users, there is another way of performing this upgrade: Change the Processor Type and Model of the existing IBM z16 processor definition to an IBM z17, and then use the Build Production feature.

3.3.1 Scenario overview

This scenario describes either the configuration steps to upgrade an existing IBM z16 A01 or an IBM z15 T01 to an IBM z17.

> **Note:** The following steps apply to an upgrade of an existing IBM z16 A01 to an IBM z17 ME1.

The key factors include:

- Hardware Configuration Definition (HCD) requires a new CPC (processor) ID for the machine type 9175 processor.
- Keep the same CPC name for the 9175 (this item is optional; the CPC name can be changed).
- The 9175 processor channels connect to the same switch ports and access the same control unit (CU) interfaces.
- The CU interfaces connect to the same switch ports.
- The starting IODF is the current 3931 *production* IODF.
- The target IODF is a new 9175 *work* IODF.
- HCD actions:
  - Migrate updated input/output configuration program (IOCP) statements.
  - Build a production IODF.
  - Remote write IODF to the input/output configuration data set (IOCDS).
- Hardware Management Console (HMC) actions:
  - Build the Reset Profile and point to the required IOCDS.
  - Build and verify the Image Profiles.
  - Build and verify the Load Profiles.
  - Perform a Power on Reset (POR).

This example uses an 3931 (IBM z16 A01) processor with a Processor ID of `PAV0` with six channel subsystems (CSSs) (CSS ID=0 - CSS ID=5). This system is replaced by a 9175 (IBM z17 ME1) with a Processor ID of `MENSA` and six CSSs. The CPC name `PAV0` and serial number are not changed.

Table 3-1 summarizes the migration options and tool requirements. The process steps are described in “HCD: Migrating the existing 3931 IODF”.

Table 3-1 I/O configuration that is migrated to a 9175 system

Options and tools	Comments
Processor ID	Must change the Processor ID to a new ID
CPC name	Local System Name. Generally should be the same name
Channel to switch port connections	Same ports
CU to switch port connections	Same ports
Starting IODF	Current active *production* IODF

Options and tools	Comments
Target IODF	Create a *work* IODF
HCD action	Repeat and change
ConnectivityCHPID Mapping Tool (CMT)	Optional, but good for verifying configurations
CFReport file (Configuration Control Number (CCN))	Required for the CMT
IOCP (Import from validated work IODF)	Yes.
CMT actions (physical channel ID (PCHID) reset)	Yes
CMT IOCP Output	Yes
CMT Reports	Yes, CHPID and CHPID to CU Report

HCD: Migrating the existing 3931 IODF

The following steps explain how to upgrade an existing 3931 processor in your IODF to the new 9175 processor by using HCD. Then, you migrate the I/O configuration and logical partitions (LPARs) from the 3931 to the 9175. Using HCD, the sequence of operations is as follows:

1. Creating the work IODF from the current 3931 production IODF.
2. Repeating the 3931 processor to be replaced.
3. Coupling Link information messages.
4. Deleting any unsupported items in the repeated 3931.
5. Add replacements for the unsupported items that were deleted in step 4.
6. Changing the 3931 to a 9175 and deleting the 3931.
7. Reconnecting the CF channel paths that were not migrated.
8. Exporting and importing OSA-ICC configuration data with OSA Advanced Facilities.
9. Using OSA Advanced Facilities to set OSA parameters.

3.3.2 Creating the work IODF from the current 3931 production IODF

HCD is the tool that is used to make a work IODF, but you start from the production IODF that contains the 3931 processor that you are upgrading (for example, `SYS9.IODF80`).

3.3.3 Repeating the 3931 processor to be replaced

To repeat the 3931 processor in HCD, complete the following steps:

1. From the main HCD panel, select option `1.3. Processor List`.
2. In the Processor List (Figure 3-1), enter `r` (for repeat) next to the 3931 that you want to upgrade, and press Enter.

```
Processor List        Row 1 of 2 More:
Command ===> ________________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
r PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-1 Processor List: Repeating processor

3. The Identify Target IODF panel opens. Do *one* of the following actions:
   - To retain all the other processor definitions in the IODF, press Enter.
   - Enter a different target IODF data set name. In this case, only the processor that you are repeating is retained in the target IODF.
4. The Create Work I/O Definition File panel prompts you to enter the data set name of the target IODF (for example, `ZNEXT01.IODF89.WORK`).
5. The Repeat Processor panel opens (Figure 3-2). Enter the Processor ID of the new 3931 (in this example, `MENSA`), keep all the other fields unchanged, and press Enter.

```
*----------------------------- Repeat Processor ------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . MENSA___                                     |
|                                                                             |
| Processor type . . . . . . . : 3931                                         |
| Processor model  . . . . . . : A01                                          |
| Configuration mode . . . . . : LPAR                                         |
|                                                                             |
| Serial number  . . . . . . . . 071A083931                                   |
| Description  . . . . . . . . . Pavo                                         |
|                                                                             |
| Specify SNA address only if part of a processor cluster:                    |
|                                                                             |
| Network name . . . . . . . . . IBM390PS  +                                  |
| CPC name . . . . . . . . . . . PAVO      +                                  |
|                                                                             |
| Local system name  . . . . . . PAVO                                         |
|                                                                             |
*--------------------------------------*                                      |
| New IODF ZNEXT01.IODF89.WORK defined. |                                     |
*--------------------------------------* ------------------------------------*
```

Figure 3-2 Repeat Processor: Defining a new Processor ID

3.3.4 Coupling Link information messages

You might receive severity messages (E, I, or W). As shown in Figure 3-3 (`CBDG441I`), severity I messages are displayed because the Coupling Facility (CF) Link CHPIDs were not copied to the new 3931 definition.

```
*------------------------------ Message List -------------------------------*
|   Save  Query  Help                                                       |
| ------------------------------------------------------------------------- |
|                                                               Row 1 of 31 |
| Command ===> ___________________________________________ Scroll ===> CSR  |
|                                                                           |
| Messages are sorted by severity. Select one or more, then press Enter.    |
|                                                                           |
| / Sev Msg. ID  Message Text                                               |
| _ I   CBDG441I The coupling facility connection between channel path      |
| #              3.80 of processor PAVO and channel path 2.D2 of            |
| #              processor VELA is not copied.                              |
| _ I   CBDG441I The coupling facility connection between channel path      |
| #              3.84 of processor PAVO and channel path 2.D6 of            |
| #              processor VELA is not copied.                              |
| _ I   CBDG441I The coupling facility connection between channel path      |
| #              3.F8 of processor PAVO and channel path 3.F9 of            |
| #              processor PAVO is not copied.                              |
| _ I   CBDG441I The coupling facility connection between channel path      |
| #              3.F9 of processor PAVO and channel path 3.F8 of            |
| #              processor PAVO is not copied.                              |
| _ I   CBDG441I The coupling facility connection between channel path      |
| #              3.FA of processor PAVO and channel path 3.FB of            |
*---------------------------------------------------------------------------*
```

Figure 3-3 Message List: Showing CBDG441I

To resolve this issue, complete the following steps:

1. Scroll until you reach the end of the messages and see the `CBDG271I Requested action on object PAVO successfully processed` message.
2. Press PF3 or PF12 to continue. As shown in Figure 3-4, there is an extra 3931 processor that is named `MENSA`.

```
Processor List        Row 1 of 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    3931     A01      LPAR  071A083931 Pavo
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-4 Processor List: Repeated processor

3.3.5 Deleting any unsupported items in the repeated 3931

If you are upgrading a processor to an IBM z17 that contains CHPID type OSE, then this CHPID type must be changed to OSD, OSH (TCP/IP) or Enterprise Extender (for SNA protocol only) configuration.

Note: CHPIDs of type OSM must be deleted in an IODF for machine type 9175.

To delete unsupported CHPIDs, complete the following steps:

1. From the Processor List panel, select the newly created `MENSA` processor and then press Enter, as shown in Figure 3-5.

```
Processor List        Row 1 of 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
s MENSA    3931     A01      LPAR  071A083931 Pavo
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-5 Processor List: Selected processor

2. On the Channel Subsystem List panel, select definitions in `CSS ID 0`, as shown in Figure 3-6.

```
Channel Subsystem List    Row 1 of 6 More:
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel subsystems, then press Enter.  To add, use F11.

Processor ID . . . : MENSA        Pavo

  CSS Devices in SS0     Devices in SS1     Devices in SS2     Devices in SS3
/ ID  Maximum + Actual  Maximum + Actual  Maximum + Actual  Maximum + Actual
s 0   65280     5806    65535     0       65535     0       65535     0
_ 1   65280     6766    65535     0       65535     0       65535     0
_ 2   65280     7070    65535     0       65535     0       65535     0
_ 3   65280     5120    65535     0       65535     0       65535     0
_ 4   65280     4126    65535     0       65535     0       65535     0
_ 5   65280     0       65535     0       65535     0       65535     0
```

Figure 3-6 Channel Subsystem List: Selected CSS

3. Within the selected CSS, set a filter, as shown in Figure 3-7.

```
 Goto  Filter  Backup  Query  Help
 ----- *-------------------------------------* ------------------------------
       | 1  1. Set Filter                    | t      Row 1 of 43 More:
 Comma |    2. Clear Filter                  | ______________ Scroll ===> CSR
       |    3. Count rows on (filtered) list |
 Selec *-------------------------------------* nter. To add use F11.

 Processor ID . . . . : MENSA        Pavo
 Configuration mode . : LPAR
 Channel Subsystem ID : 0

         CHID+            Dyn Entry +
 / CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
 _ 10    100   FC    SPAN  01  01 1C       No  TS7760A
 _ 11    144   FC    SPAN  01  01 1D       No  TS7760A
 _ 12    184   FC    SPAN  02  02 1C       No  TS7760A
 _ 13    1C0   FC    SPAN  02  02 1D       No  TS7760A
 _ 40    101   FC    SPAN  01  01 24       No  DASD
 _ 41    145   FC    SPAN  02  02 24       No  DASD
 _ 42    185   FC    SPAN  01  01 25       No  DASD
 _ 43    1C1   FC    SPAN  02  02 25       No  DASD
 _ 44    120   FC    SPAN  01  01 1E       No  FCTC
 _ 45    160   FC    SPAN  02  02 1E       No  FCTC
 _ 50    121   FC    SPAN  01  01 2A       No  DASD
```

Figure 3-7 Channel Path List: Set Filter

4. In the Filter Channel Path List panel, select the channel path type OSE, as shown in Figure 3-8.

```
*----------------------- Filter Channel Path List ------------------------*
|                                                                          |
|                                                                          |
| Specify or revise the following filter criteria.                         |
|                                                                          |
| Channel path type  . OSE                                                 |
| Operation mode . . . ____  +                                             |
| Managed  . . . . . . _    (Y = Yes; N = No)   I/O Cluster ________   +  |
| Dynamic entry switch __    +                                             |
| Entry switch . . . . __    +                                             |
| CF connected . . . . _    (Y = Connected; N = Not connected)            |
| CHID AID/P PCHID/P   _____                                               |
|                                                                          |
| Description  . . . . ________________________________                    |
|                                                                          |
| Partition  . . . . . ________    +                                       |
| Connected to CUs . . _    (Y = Connected; N = Not connected)            |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
*--------------------------------------------------------------------------*
```

Figure 3-8 Filter Channel Path type: Type OSE for channel path type

5. The Channel Path List panel shows channel definitions only for channel path type OSE. Delete these definitions as shown in Figure 3-9 on page 48, and press Enter.

```
                              Channel Path List     Filter Mode. More:
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Pavo
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
d B4    1CC   OSE   SPAN  __  __ __       No  1KBT  7S On loan to Redbooks
```

Figure 3-9 Channel path list: Delete all channel definitions for type OSE

6. Confirm that spanned channels are deleted from all accessing CSSs by pressing Enter, as shown in Figure 3-10.

```
*--------------------- Confirm Delete Channel Path ---------------------*
|                                                           Row 1 of 1 | --
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Scroll forward to view the complete list of channel paths to be       |
| deleted. Press ENTER to confirm delete request. Press F12 to cancel   |
| delete request.                                                       |
|                                                                       |
| Processor ID . . . . : MENSA        Pavo                              |
| Channel Subsystem ID : 0                                              |
|                                                                       |
|                    CHID                                               |
| CHPID   Type  Mode  AID/P                                             |
| B4      OSE   SPAN  1CC                                               |
| ************************** Bottom of data *************************** | **
|                                                                       |
|                                                                       |
|                                                                       |
|                                                                       |
|                                                                       |
|                                                                       |
*-----------------------------------------------------------------------------*
| Spanned channel paths are deleted from all accessing channel subsystems. |
*-----------------------------------------------------------------------------*
```

Figure 3-10 Confirm Delete Channel Path panel

Note: When upgrading a processor that contains any PCIe FUNCTION statements of type ROCE or ROCE-2, these FUNCTION statements must be deleted from the IODF before changing the processor type to 9175. These FUNCTION statements have been replaced by NETH FUNCTION statements.

See Chapter 15 for the new NETH PCIe function that supports SMC-R for IBM z17 Chapter 15, “Adding Peripheral Component Interconnect Express devices” on page 331.

7. In case there are ROCE-2 `FUNCTIONS` defined on the 3931 processor, they have to be removed. If the equivalent `FUNCTIONS` have to be defined again on the 9175, a printout of the current definitions for reference can be produced. From the HCD main panel select option three Print or compare configuration data, see Figure 3-11

```
z/OS V3.1 HCD
Command ===>
_________________________________________________________________

                          Hardware Configuration

Select one of the following.

3   0.  Edit profile options and policies
    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 . . . 'ZNEXT01.IODF89.WORK'                  +
```

Figure 3-11 Select Print option

8. From the Print or Compare Configuration Data panel select option 1 Print configuration reports, as shown in Figure 3-12

```
*------------ Print or Compare Configuration Data ------------*
|                                                             |
|                                                             |
| Select one of the following tasks.                          |
|                                                             |
| 1_  1.  Print configuration reports                         |
|     2.  Print the activity log                              |
|     3.  View the activity log                               |
|     4.  Compare IODFs                                       |
|     5.  Compare CSS / operating system views                |
|                                                             |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 3-12 Select Print configuration reports

9. From Print Configuration Reports panel select CSS report, as shown in figure Figure 3-13 on page 50

```
*------------------------ Print Configuration Reports -------------------------*
|                                                                              |
|                                                                              |
| Select the types of report you want, and specify the values below.          |
|                                                                              |
| IODF name  : 'ZNEXT01.IODF89.WORK'                                          |
|                                                                              |
| Types of report                  Limit report(s)                            |
| /  CSS report                    1  1. Yes                                  |
| _  Switch report                    2. No                                   |
| _  OS report                                                                |
| _  CTC connection report                                                    |
| _  I/O path report                                                          |
|    (may invoke I/O Autoconfiguration (zDAC))                                |
|                                                                              |
| Job statement information                                                   |
| //ZNEXT01H  JOB (SG248581),'FPINTO HCD',                                    |
| //          MSGCLASS=X,CLASS=A,MSGLEVEL=(1,1),                              |
| //          NOTIFY=&SYSUID                                                  |
| /*JOBPARM LINES=9999                                                        |
| //*                                                                         |
| //*                                                                         |
|                                                                              |
*------------------------------------------------------------------------------*
```

Figure 3-13 Select CSS report to save Functions information

10. From the Available CSS Report Types option panel select CSS summary reports

```
*------- Available CSS Report Types --------*
|                                           |
|                                           |
| Select one or more.                       |
|                                           |
| /  CSS summary reports                    |
| _  Channel path detail reports            |
| _  Control unit detail report             |
| _  Device detail report                   |
|                                           |
|                                           |
|                                           |
|                                           |
*-------------------------------------------*
```

Figure 3-14 Select CSS summary reports from the CSS Report Types options panel.

11. From the Limit Reports panel select the `MENSA` processor ID

```
*------------------------ Limit Reports -------------------------*
|                                                                |
|                                                                |
| To limit the reports, specify the following criteria related   |
| to the IODF in access.                                         |
|                                                                |
|                                                                |
|                                 Applicable for:                |
| Processor ID . . . . MENSA    + CSS, CTC, I/O path reports     |
| Partition name . . . ________ + CSS, CTC, I/O path report      |
| OS configuration ID  ________ + OS, I/O path report            |
| Switch ID  . . . . . __       + switch report                  |
|                                                                |
| Specify the sysplex and system name to gather the actual       |
| configuration from. (Blanks default to the local system.)      |
|                                                                |
| Sysplex name . . . . ________   I/O path report                |
| System name  . . . . ________   I/O path report                |
|                                                                |
|                                                                |
|                                                                |
*----------------------------------------------------------------*
```

Figure 3-15 Select processor ID of the previously defined processor

12. After the resulting Batch job has successfully run, a list of defined `FUNCTION` statements is printed, see Figure 3-16.

```
 PCIE FUNCTION SUMMARY REPORT
 PROCESSOR ID  MENSA    TYPE  3931     MODEL  A01      CONFIGURATION MODE: LPAR

FID  VF   PF CHID  PORT  TYPE         PM UID   DESCRIPTION
____ ___  __ _____ ____  ____________ __ ____  ________________________________
2435   5     19C   2     ZHYPERLINK
3101   1     11C   1     ROCE-2                10GbE
3131   2     11C   1     ROCE-2                10GbE
3132   3     11C   1     ROCE-2                10GbE
3133   4     11C   1     ROCE-2                10GbE
3134   5     11C   1     ROCE-2                10GbE
3135   6     11C   1     ROCE-2                10GbE
3138  11     11C   1     ROCE-2                10GbE
3145   7     11C   1     ROCE-2                10GbE
3146   8     11C   1     ROCE-2                10GbE
3147   9     11C   1     ROCE-2                10GbE
3148  10     11C   1     ROCE-2                10GbE
3201   1     11C   2     ROCE-2                10GbE
3231   2     11C   2     ROCE-2                10GbE
3232   3     11C   2     ROCE-2                10GbE
3233   4     11C   2     ROCE-2                10GbE
3234   5     11C   2     ROCE-2                10GbE
3235   6     11C   2     ROCE-2                10GbE
```

Figure 3-16 Partial output of CSS Report

13. To delete the ROCE-2 Functions select option f for `MENSA` on the Processor List panel, see Figure 3-17 on page 52

```
 Processor List        Row 1 of 3 More:
 Command ===> _______________________________________________ Scroll ===>
CSR

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


 / Proc. ID Type +   Model +  Mode+ Serial-# + Description
 f MENSA    3931     A01      LPAR  071A083931 Pavo
 _ PAVO     3931     A01      LPAR  071A083931 Pavo
 _ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-17 Select MENSA to work with PCIe FUNCTIONS

14.On the PCIe Function List panel set a filter for ROCE-2, see Figure 3-18

```
 Goto  Filter  Backup  Query  Help
 ----- *-------------------------------------*
------------------------------
       | 1  1. Set Filter                    | st    Row 1 of 136 More:
 Comma |    2. Clear Filter                  | ______________ Scroll ===>
CSR
       |    3. Count rows on (filtered) list |
 Selec *-------------------------------------* Enter. To add, use F11.

 Processor ID . . . . : MENSA       Pavo



 / FID    CHID+ P+  VF+ PF Type+        UID   Description
 _ 1031   7C6   _   1   _  ISM          ____  SC76
 _ 1032   7C6   _   2   _  ISM          ____  SC74
 _ 1033   7C6   _   3   _  ISM          ____  SC75
 _ 1034   7C6   _   4   _  ISM          ____  SC80
 _ 1035   7C6   _   5   _  ISM          ____  SC81
 _ 1036   7C6   _   6   _  ISM          ____  RDBKZVM0
 _ 1037   7C6   _   7   _  ISM          ____  RDBKESI3
 _ 1038   7C6   _   8   _  ISM          ____  SC01
 _ 1039   7C6   _   9   _  ISM          ____  WTSCNET
 _ 103A   7C6   _   10  _  ISM          ____  WTSCMXA
 _ 103B   7C6   _   11  _  ISM          ____  RDBKPOK
 _ 1041   7C6   _   12  _  ISM          ____  RDBKZVM1
```

Figure 3-18 Set a filter for a specific FUNCTION

15.Mark all the resulting entries and select delete on the following panel

```
*----------------------- Filter PCIe Function List -----------------------*
|                                                                         |
|                                                                         |
| Specify or revise the following filter criteria.                        |
|                                                                         |
| CHID  . . . . . . . . . ___                                             |
| PORT  . . . . . . . . . _                                               |
|                                                                         |
| Type  . . . . . . . . . ROCE-2       +                                  |
|                                                                         |
| Virtual Function ID . . ___                                             |
|                                                                         |
| Description . . . . . . ________________________________                |
|                                                                         |
| Partition . . . . . . . ________  +                                     |
|                                                                         |
| UIDs defined  . . . . . _  (Y = YES; N = NO)                            |
| UID . . . . . . . . . . ____                                            |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 3-19 Select ROCE-2 as the PCIe function list type and press ENTER

16. On the resulting panel, mark all ROCE-2 functions for deletion and select delete on the pop up panel, see Figure 3-20

```
Goto  Filter  Backup  Query  Help
------------- *------------- Actions on selected PCIe Function -------------*
              |                                                             |
Command ===>  |                                                             |
              | Select by number or action code and press Enter.            |
Select one or |                                                             |
              |d__  1.  Add like . . . . . . . . . . . . . . . (a)          |
Processor ID  |     2.  Change . . . . . . . . . . . . . . . . (c)          |
              |     3.  Delete . . . . . . . . . . . . . . . . (d)          |
              |     4.  View . . . . . . . . . . . . . . . . . (v)          |
              |                                                             |
/ FID   CHID  |                                                             |
( 3101  11C   *-------------------------------------------------------------*
_ 3131  11C   1   2   _  ROCE-2        ____  10GbE
_ 3132  11C   1   3   _  ROCE-2        ____  10GbE
_ 3133  11C   1   4   _  ROCE-2        ____  10GbE
_ 3134  11C   1   5   _  ROCE-2        ____  10GbE
_ 3135  11C   1   6   _  ROCE-2        ____  10GbE
_ 3138  11C   1   11  _  ROCE-2        ____  10GbE
_ 3145  11C   1   7   _  ROCE-2        ____  10GbE
_ 3146  11C   1   8   _  ROCE-2        ____  10GbE
_ 3147  11C   1   9   _  ROCE-2        ____  10GbE
_ 3148  11C   1   10  _  ROCE-2        ____  10GbE
) 3201  11C   2   1   _  ROCE-2        ____  10GbE
```

Figure 3-20 Use “(” and “)” to mark a block of entries for deletion

3.3.6 Changing the 3931 to a 9175 and deleting the 3931

You can either keep the original copy of the 3931 (PAV0) or delete it from the IODF. In this example, keep it in the IODF for a few more steps.

To change the 3931 to a 9175, complete the following steps:

1. Enter c (for change) next to PAV0 to change the 3931 to a 9175 and press Enter. The Change Process Definition panel opens (Figure 3-21).
2. Make the following updates, and press Enter:
   - Update Processor type to 9175.
   - Update Processor model to ME1.
   - Update the 3931 part of the Serial number to 9175 (that is, 012345*3931* to 012345*9175*).
   - Update Description to Mensa.
   - Update Local system name to Mensa.

> **Note:** Because in this example we use HCD to write an IOCDS to this 3931 in preparation for an upgrade, we must leave the Network name and CPC name set to IBM390PS and PAV0. These settings must be updated in the IODF after the 3931 is upgraded to a 9175.

```
*------------------------ Change Processor Definition ------------------------*
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Processor ID . . . . . . . . : MENSA                                         |
| Support level:                                                               |
| 3931 support                                                                 |
| Processor type . . . . . . . . 9175     +                                    |
| Processor model  . . . . . . . ME1      +                                    |
| Configuration mode . . . . . . LPAR     +                                    |
|                                                                              |
| Serial number  . . . . . . . . 071A089175  +                                 |
| Description  . . . . . . . . . Mensa                                         |
|                                                                              |
| Specify SNA address only if part of a processor cluster:                     |
|                                                                              |
| Network name . . . . . . . . . IBM390PS  +                                   |
| CPC name . . . . . . . . . . . PAV0     +                                    |
|                                                                              |
| Local system name  . . . . . . Mensa                                         |
|                                                                              |
|                                                                              |
*------------------------------------------------------------------------------*
```

Figure 3-21 Processors: Change Processor Definition

3. The Update Channel Path Identifiers panel opens (Figure 3-22). No changes are made in this example.

```
*-------------------- Update Channel Path Identifiers ----------------------*
|                                                              Row 1 of 42 |
| Command ===> __________________________________________ Scroll ===> CSR   |
|                                                                           |
| Specify any changes to the channel path identifiers in the list below.    |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Channel Subsystem ID : 0                                                  |
|                                                                           |
| CHPID  Type  Side  Until CHPID   New CHPID +                              |
| 10     FC          __            10                                       |
| 11     FC          __            11                                       |
| 12     FC          __            12                                       |
| 13     FC          __            13                                       |
| 40     FC          __            40                                       |
| 41     FC          __            41                                       |
| 42     FC          __            42                                       |
| 43     FC          __            43                                       |
| 44     FC          __            44                                       |
| 45     FC          __            45                                       |
| 50     FC          __            50                                       |
| 51     FC          __            51                                       |
| 52     FC          __            52                                       |
*---------------------------------------------------------------------------*
```

Figure 3-22 Processors: Update Channel Path Identifiers

4. Press Enter for each CSS ID.
5. The repeated 3931 processor is successfully changed to a 9175-ME1, as shown in Figure 3-23.

```
Processor List          Row 1 of 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  071A089175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-23 Processor List: Changed processor

3.3.7 Deleting the 3931 processor definition

Now that the 3931 is repeated and changed to a 9175, the original 3931 definition (PAVO) must be deleted so that the required CF links can be restored.

To delete the 3931 processor definition, complete the following steps:

1. Enter `d` (for delete) next to the `PAVO` processor in the Processor List panel (Figure 3-24).

```
Processor List        Row 1 of 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  071A089175 Mensa
d PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-24 Processor List: Deleting the processor

2. Press Enter to confirm the deletion of the processor (Figure 3-25).

```
*--------------------- Confirm Delete Processor ----------------------*
|                                                          Row 1 of 1 |
| Command ===> _______________________________________ Scroll ===> CSR |
|                                                                      |
| Scroll forward to view the complete list of processors to be        |
| deleted. Press ENTER to confirm delete request. Press F12 to        |
| cancel delete request.                                               |
|                                                                      |
| Processor ID Type     Model    Description                          |
| PAVO         3931     A01      Pavo                                 |
| ************************ Bottom of data ************************** |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
|                                                                      |
*---------------------------------------------------------------------*
```

Figure 3-25 Confirm delete of source processor

3. Source processor 3931 has been deleted.

```
Processor List        Row 1 of 2 More:
Command ===> ______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  071A089175 Mensa
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-26 Processor List: Processor deleted

3.3.8 Reconnecting the CF channel paths that were not migrated

Manually redefine the CF Links that you want from the `MENSA` (previously `PAVO`) processor to any other processor, along with any internal CF links that you want. To help in this effort, you can get a CF connection report from the previous production IODF containing the 3931. Alternatively, you can make a note of all `CBDG441I` error messages that you received in 3.3.7, “Deleting the 3931 processor definition” on page 55.

3.3.9 Additional steps and tasks

When you are ready to map the PCHIDs from the IBM z17 CFReport file to the CHPIDs in your exported IODF, go to Chapter 4, “Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool” on page 63.

To define the I/O configuration for your system, go to Chapter 5, “Building the production input/output definition file and setting up the central processor complex” on page 89.Installing a new IBM z17 system into an existing IBM Z environment

This section describes the steps for adding an IBM z17 ME1 into an existing IBM Z environment.

3.4 Installing a new IBM z17 scenario overview

This scenario shows the configuration steps for defining a new 9175 processor in an existing hardware environment.

The process has the following key considerations:

- HCD requires a new processor ID for the 9175.
- HCD requires a new CPC name for the 9175.
- The 9175 processor connects to new switch ports and new CU interfaces.
- The CU interfaces connect to the same switch ports as they did previously.
- The starting IODF is the current production IODF.
- The target IODF is a new work IODF with a 9175 defined.
- HCD actions: Migrate updated IOCP statements.
  - Build a production IODF.
  - Remote write an IODF to IOCDS.
- The HMC actions:
  - Build Reset Profile and point to required IOCDS.
  - Build and verify Image Profiles.
  - Build and verify Load Profiles.

– Run a POR.

This example defines a new 9175 (IBM z17 ME1) processor with a Processor ID of `MENSA2` and with six CSSs (CSS ID=0 - CSS ID=5). The CPC name of `MENSA2` and serial number of 02-B9FB8 are used for the 9175.

Table 3-2 summarizes the tool requirements.

Table 3-2 I/O configuration for a new (additional) 9175 processor

New (additional) 9175 processor	New (additional) 9175 processor to connect to the new switch ports and same CUs to which existing processors connect
Processor ID	Requires a new Processor ID.
CPC name	Requires a new CPC name.
Channel to switch port connections	Extra ports.
CU to switch port connections	Same ports.
Starting IODF	Current active production IODF.
Target IODF	Create a work IODF.
HCD action	Add processor.
CMT Program	Optional, but good for verifying configurations.
CFReport File (CCN)	Required for the CMT.
IOCP (import from validated work IODF)	Yes.
CMT actions (PCHID reset)	Yes.
CMT IOCP Output	Yes.
CMT Reports	Yes, CHIPID Report and CHIPID to CU Report.

HCD: Creating a 9175 IODF

The following steps explain how to define an additional 9175 processor in your existing IODF to the existing I/O configuration by using HCD:

1. Creating a work IODF from the current production IODF.
2. Adding the 9175 processor.
3. Exporting and importing OSA-ICC configuration data with OSA Advanced Facilities.
4. Using OSA Advanced Facilities to set OSA parameters.

3.4.1 Creating a work IODF from the current production IODF

HCD is the tool that is used to make a work IODF. In this example, we start from the current production IODF that contains the existing hardware environment that will be connected to the new 9175 processor (for example, `SYS9.IODF80`).

3.4.2 Adding the 9175 processor

To add the 9175 processor, complete the following steps:

1. From the HCD main menu, select option `1.3. Processor List`.
2. In the Processor List (Figure 3-27), press PF11, or enter `add` on the command line to add a processor, and press Enter.

```
 Processor List        Row 1 of 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  0123459175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
```

Figure 3-27 Processor List: Adding a processor

The Add Processor panel opens (Figure 3-28).

```
*------------------------------- Add Processor -------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . ________                                     |
| Processor type . . . . . . . . ________  +                                  |
| Processor model  . . . . . . . ________  +                                  |
| Configuration mode . . . . . . LPAR      +                                  |
| Number of channel subsystems . . _       +                                  |
|                                                                             |
| Serial number  . . . . . . . . __________                                   |
| Description  . . . . . . . . . ________________________________             |
|                                                                             |
| Specify SNA address only if part of a processor cluster:                    |
|                                                                             |
| Network name . . . . . . . . . ________  +                                  |
| CPC name . . . . . . . . . . . ________  +                                  |
|                                                                             |
| Local system name  . . . . . . ________                                     |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 3-28 Add Processor: Data fields to be updated

3. Specify the appropriate values. For example, specify the following settings, as shown in Figure 3-29 on page 61:

Processor ID	`MENSA2`
Processor type	`9175`
Processor model	`ME1`
Number of channel subsystems	(Keep this blank for now.)
Serial number	`02B9FB89175`
Network name	`IBM390PS`
CPC name	`MENSA2`
Local System Name	(Keep this blank for now.)

```
*------------------------------- Add Processor -------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . . MENSA2                                     |
| Processor type . . . . . . . . . 9175      +                                |
| Processor model  . . . . . . . . ME1       +                                |
| Configuration mode . . . . . . . LPAR      +                                |
| Number of channel subsystems . . _         +                                |
|                                                                             |
| Serial number  . . . . . . . . . 2B9FB89175                                 |
| Description  . . . . . . . . . . ________________________________           |
|                                                                             |
| Specify SNA address only if part of a processor cluster:                    |
|                                                                             |
| Network name . . . . . . . . . . IBM390PS  +                                |
| CPC name . . . . . . . . . . . . MENSA2    +                                |
|                                                                             |
| Local system name  . . . . . . . ________                                   |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 3-29 Add Processor: Data fields updated

4. Press Enter. The Create Work I/O Definition File panel opens and prompts you to enter the data set name of the target IODF (for example, `SYS9.IODF90.WORK`).
5. Press Enter. You now have a 9175 processor that is named `MENSA2` (Figure 3-30).

```
 Processor List        Row 1 of 4 More:
 Command ===> _______________________________________________ Scroll ===> CSR

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


 / Proc. ID Type +   Model +  Mode+ Serial-# + Description
 _ MENSA    9175     ME1      LPAR  0123459175 Mensa
 _ MENSA2   9175     ME1      LPAR  2B9FB89175 ______________________________
 _ PAVO     3931     A01      LPAR  071A083931 Pavo
 _ VELA     3932     AGZ      LPAR  087F283932 Vela
 ******************************* Bottom of data *****************************




*------------------------------------------------------------------*
| Definition of processor MENSA2 has been extended to its maximum |
| configuration.                                                   |
*------------------------------------------------------------------*
```

Figure 3-30 Processor List: New processor added

The message at the bottom of the panel is generated because the HCD automatically populated the processor with all allowed CSSs and reserved partitions. In HCD, when you define a new or redefine a processor as a 9175, HCD no longer defines or allow you to define partitions 0xB to 0xF in CSS5. These partitions are reserved for IBM internal use. HCD automatically defines the maximum configuration of 6 CSSs and 85 LPARs.

6. Enter `s` next to `MENSA2`, and press Enter. The Channel Subsystem List panel opens. Here you can see six CSSs (CSS0 - CSS5) that are defined with the default `MAXDEV` values for `SS0` of `65280` set by HCD and `65535` set for `SS1`, `SS2`, and `SS3` (Figure 3-31).

```
Channel Subsystem List    Row 1 of 6 More:
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel subsystems, then press Enter.  To add, use F11.

Processor ID . . . : MENSA2

  CSS Devices in SS0    Devices in SS1    Devices in SS2    Devices in SS3
/ ID  Maximum + Actual  Maximum + Actual  Maximum + Actual  Maximum + Actual
_ 0   65280     0       65535     0       65535     0       65535     0
_ 1   65280     0       65535     0       65535     0       65535     0
_ 2   65280     0       65535     0       65535     0       65535     0
_ 3   65280     0       65535     0       65535     0       65535     0
_ 4   65280     0       65535     0       65535     0       65535     0
_ 5   65280     0       65535     0       65535     0       65535     0
```

Figure 3-31 Channel Subsystem List: Four Subchannel Sets

3.4.3 Additional steps and tasks

When you are ready to map the PCHIDs from the IBM z17 CFReport file to the CHPIDs in your exported IODF, go to Chapter 4, “Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool” on page 63.

To define the I/O configuration for your system, go to Chapter 5, “Building the production input/output definition file and setting up the central processor complex” on page 89.

4

Preparing an IOCP input file to use the IBM Z Connectivity Mapping Tool

This chapter describes in detail how to use the Connectivity Mapping Tool (CMT).

Naming: The IBM z system that is targeted by this publication is IBM z17 ME1. Throughout this chapter, we might refer to this machine as IBM z17.

We will describe a series of steps where a new IBM z17 will be installed in place of an already running IBM z16.

This chapter includes the following topics:

- Creating a copy of an existing IODF from an IBM z16 server
- Creating the IOCP file for use with the CMT
- Assigning CHIDs to CHPIDs by using the CMT
- Importing the CFR file and the IOCP input file into CMT
- Building the initial Overview
- Assigning missing feature types
- Resetting Incompatible (Hardware - I/O) Entries
- Creating CMT reports
- Creating an updated IOCP file
- Additional steps and processes

4.1 Creating a copy of an existing IODF from an IBM z16 server

To create an exact copy of a defined IBM Z server in HCD go to the Main HCD Panel, select the current IODF file. Determine the currently active IODF file by using a command processor and entering command `D IOS,CONFIG` as shown in the Example 4-1.

Example 4-1 Display active IODF filename

```
D IOS CONFIG
IOS506I 14.17.31 I/O CONFIG DATA 727
ACTIVE IODF DATA SET = SYS9.IODF91
CONFIGURATION ID = ITSO          EDT ID = 01
TOKEN:  PROCESSOR DATE     TIME     DESCRIPTION
 SOURCE: PAVO     25-05-02 14:49:52 SYS9     IODF91
ACTIVE CSS:  3    SUBCHANNEL SETS CONFIGURED: 0, 1, 2, 3
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE
SUBCHANNEL SET FOR PPRC PRIMARY: INITIAL = 0   ACTIVE = 0
HYPERSWAP FAILOVER HAS OCCURRED: NO
LOCAL SYSTEM NAME (LSYSTEM): PAVO
```

Next, start the HCD application and enter the active IODF data set name into the I/O definition file space, enter option "1" for `Define, modify or view configuration data` as shown in Figure 4-1.

```
z/OS V3.1 HCD
Command ===> ________________________________________________________________

                           Hardware Configuration

Select one of the following.

1_  0.  Edit profile options and policies
    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 . . . 'SYS9.IODF91'                          +
```

Figure 4-1 Initialize HCD with the active IODF file

Select option `3. Processors` to access the `Processor List` panel and add the active IODF file.

Select the source processor for the copy operation by placing an "r" in front of the CPC to be copied.

```
Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                                Processor List       Row 1 of 3 More:       >
Command ===> ______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  0B9FB89175 Mensa
r PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
******************************* Bottom of data ********************************
```

Figure 4-2   Select the processor to be copied

The `Identify Target IODF` panel appears where you can insert the target IODF into which the configuration data will be repeated. We chose to leave it as it is.

An attempt to make a change in an active IODF automatically leads to a copy operation of the selected active IODF. Therefore in the next panel `Create Work I/O Definition File`, specify a new target IODF name for the work file. Adjust the space allocation of the new file according to the amount of changes you plan to apply. Please adhere to the valid rules for allocating new datasets in your environment, see Figure 4-3.

```
Goto  Filter  Backup  Query  Help
- *------------------------ Identify Target IODF --------------------------*
  |                                                                        |
C |   *----------------- Create Work I/O Definition File -----------------* |
  | S |                                                                   | |
S |   |                                                                   | |
  | T | The current IODF is a production IODF and therefore cannot be     | |
  |   | updated. To create a new work IODF based on the current           | |
/ |   | production IODF, specify the following values.                    | |
_ |   |                                                                   | |
r |   | IODF name  . . . . . . 'ZNEXT01.IODF92.WORK'                       | |
_ *-- |                                                                   | *
***** | Volume serial number . BH5ST6  +                                  | ***
      |                                                                   |
      | Space allocation . . . 2048    (Number of 4K blocks)              |
      | Activity logging . . . no     (Yes or No)                         |
      | Multi-user access  . . No     (Yes or No)                         |
      |                                                                   |
      |                                                                   |
      *-------------------------------------------------------------------*
```

Figure 4-3   Specify the allocation parameters for the work IODF file

On the `Repeat Processor` panel enter the values for `Processor ID`, `Description` and `Local System Name`, see Figure 4-4 on page 66.

```
*---------------------------- Repeat Processor ------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . PAVO2___                                     |
|                                                                             |
| Processor type . . . . . . . : 3931                                         |
| Processor model  . . . . . . : A01                                          |
| Configuration mode . . . . . : LPAR                                         |
|                                                                             |
| Serial number  . . . . . . . . 071A083931                                   |
| Description  . . . . . . . . . Pavo2                                        |
|                                                                             |
| Specify SNA address only if part of a processor cluster:                    |
|                                                                             |
| Network name . . . . . . . . . IBM390PS  +                                  |
| CPC name . . . . . . . . . . . PAVO      +                                  |
|                                                                             |
| Local system name  . . . . . . PAVO2                                        |
|                                                                             |
*-----------------*                                                           |
| Input required. |                                                           |
*-----------------* ----------------------------------------------------------*
```

Figure 4-4 Specify the new processor ID, description and the local system name

The copy operation will not copy definitions that will lead to duplicate entries. Therefore be prepared to see a number of informational messages, see Example 4-2

Example 4-2 Messages that may appear during the copy operation

```
/ Sev Msg. ID  Message Text
_ I   CBDG441I The coupling facility connection between channel path
#              3.80 of processor PAVO and channel path 2.D2 of
#              processor VELA is not copied.
_ I   CBDG441I The coupling facility connection between channel path
#              3.81 of processor PAVO and channel path 3.86 of
#              processor MENSA is not copied.
_ I   CBDG441I The coupling facility connection between channel path
#              3.84 of processor PAVO and channel path 2.D6 of
#              processor VELA is not copied.
_ I   CBDG441I The coupling facility connection between channel path
#              3.85 of processor PAVO and channel path 3.84 of
#              processor MENSA is not copied.
_ I   CBDG441I The coupling facility connection between channel path
#              3.88 of processor PAVO and channel path 3.88 of
_ I   CBDG271I Requested action on object PAVO successfully processed.
```

Make sure the copy operation was processed successfully. If this is the case, an additional Processor is created, including all basic definitions of the source processor, see Figure 4-5.

```
Goto  Filter  Backup  Query  Help
------------------------------------------------------------------------------
                                Processor List       Row 1 of 4 More:       >
Command ===> ______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ MENSA    9175     ME1      LPAR  0B9FB89175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ PAVO2    3931     A01      LPAR  071A083931 Pavo2
_ VELA     3932     AGZ      LPAR  087F283932 Vela
******************************* Bottom of data ********************************
```

Figure 4-5 Note the newly created processor

4.1.1 Change the Processor Type

Now that the selected processor has been copied, the processor type needs to be changed. On the `Processor List` panel, select the processor by inserting a “`c`” in front to change the machine type. On the `Change Processor Definition` panel, apply the change for the appropriate processor type, model and serial number. Use the “F4” key for prompting.

```
*------------------------ Change Processor Definition ------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . : PAVO2                                        |
| Support level:                                                              |
| 3931 support                                                                |
| Processor type . . . . . . . . 9175      +                                  |
| Processor model  . . . . . . . ME1       +                                  |
| Configuration mode . . . . . . LPAR      +                                  |
|                                                                             |
| Serial number  . . . . . . . . 071A089175  +                                |
| Description  . . . . . . . . . Pavo2                                        |
|                                                                             |
| Specify SNA address only if part of a processor cluster:                    |
|                                                                             |
| Network name . . . . . . . . . IBM390PS  +                                  |
| CPC name . . . . . . . . . . . PAVO      +                                  |
|                                                                             |
| Local system name  . . . . . . PAVO2                                        |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 4-6 Change processor type and serial number

The change process will not allow unsupported features to be carried over to the new processor type.

```
*------------------------------- Message List --------------------------------*
|   Save  Query  Help                                                         |
| --------------------------------------------------------------------------- |
|                                                                Row 1 of 184 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Messages are sorted by severity. Select one or more, then press Enter.      |
|                                                                             |
| / Sev Msg. ID  Message Text                                                 |
| _ E   CBDG614I Function type ROCE-2 of function 3101 is not supported       |
| #              by processor PAVO2.                                          |
| _ E   CBDG614I Function type ROCE-2 of function 3131 is not supported       |
| #              by processor PAVO2.                                          |
| _ E   CBDG614I Function type ROCE-2 of function 3132 is not supported       |
| #              by processor PAVO2.                                          |
| _ E   CBDG614I Function type ROCE-2 of function 3133 is not supported       |
| #              by processor PAVO2.                                          |
| _ E   CBDG614I Function type ROCE-2 of function 3134 is not supported       |
| #              by processor PAVO2.                                          |
| _ E   CBDG614I Function type ROCE-2 of function 3135 is not supported       |
| #              by processor PAVO2.                                          |
*-------------------------------------------------------------* pported       |
| The change of processor rules leads to invalid definitions. |               |
*-------------------------------------------------------------* --------------*
```

Figure 4-7 Partial list of errors while trying to change the processor type

Delete any unsupported entries from the copied processor prior to change the processor type. Follow the indication of the error messages as shown in Figure 4-7.

4.1.2 Validating the work input/output definition file

To validate the work input/output definition file (IODF) by using the Hardware Configuration Definition (HCD) component, complete the following steps:

1. Select HCD option `2.12. Build validated work I/O definition file`. Review the message list and correct any errors.
2. Press PF3 to continue. The `Requested action successfully processed` message is displayed.

3. Select HCD option `6.4. View I/O Definition File Information`. The IODF type is now indicated as `Work - Validated` (see Figure 4-8).

```
*---------------- View I/O Definition File Information -----------------*
|                                                                        |
|                                                                        |
| IODF name  . . . . . . : 'ZNEXT01.IODF92.WORK'                         |
| IODF type  . . . . . . : Work - Validated                              |
| IODF version . . . . . : 5                                             |
|                                                                        |
| Creation date  . . . . : 2025-05-07                                    |
| Last update  . . . . . : 2025-05-07  15:26                             |
|                                                                        |
| Volume serial number . : BH5ST6                                        |
| Allocated space  . . . : 2048    (Number of 4K blocks)                 |
| Used space . . . . . . : 768     (Number of 4K blocks)                 |
|    thereof utilized (%)  94                                            |
| Activity logging . . . : No                                            |
| Multi-user access  . . : No                                            |
| Backup IODF name . . . :                                               |
|                                                                        |
| Description  . . . . . :                                               |
|                                                                        |
|                                                                        |
|                                                                        |
| ENTER to continue.                                                     |
*------------------------------------------------------------------------*
```

Figure 4-8 View I/O Definition File Information: Validated work IODF

4.2 Creating the IOCP file for use with the CMT

To create the input/output configuration program (IOCP) input data set for the CMT, complete the following steps:

1. Select HCD option `2.3. Build IOCP input data set`, and press Enter (see Figure 4-9).

```
*------ Activate or Process Configuration Data ------*
|                                                   |
|                                                   |
| Select one of the following tasks.                |
|                                                   |
| 3   1.  Build production I/O definition file      |
|     2.  Build IOCDS                               |
|     3.  Build IOCP input data set                 |
|     4.  Create JES3 initialization stream data    |
|     5.  View active configuration                 |
|     6.  Activate or verify configuration          |
|         dynamically                               |
|     7.  Activate configuration sysplex-wide       |
|     8.  (no longer supported)                     |
|     9.  (no longer supported)                     |
|     10. Build I/O configuration data              |
|     11. Build and manage processor cluster        |
|         IOCDSs, IPL attributes and dynamic I/O    |
|         changes                                   |
|     12. Build validated work I/O definition file  |
|                                                   |
|                                                   |
|                                                   |
*---------------------------------------------------*
```

Figure 4-9 Activate or Process Configuration Data: Building IOCP for PAVO2

2. HCD displays the list of available processors (see Figure 4-10). Select the `PAVO2` processor by entering a (`s`) next to it and pressing Enter.

```
*------ Activate or Process Configuration Data ------*
C |                                                   | ___________________
  *------------------------- Available Processors ---------------------------*
  |                                                               Row 1 of 4 |
  | Command ===> __________________________________________________________  |
S |                                                                          |
  | Select one.                                                              |
2 |                                                                          |
  |   Processor ID  Type     Model    Mode  Description                      |
  |   MENSA         9175     ME1      LPAR  Mensa                            |
  |   PAVO          3931     A01      LPAR  Pavo                             |
  | s PAVO2         9175     ME1      LPAR  Pavo2                            |
  |   VELA          3932     AGZ      LPAR  Vela                             |
  | *************************** Bottom of data ***************************** |
  |                                                                          |
  |                                                                          |
  |                                                                          |
  |                                                                          |
F |                                                                          |
  |                                                                          |
I |                                                                          |
  |                                                                          |
  |                                                                          |
  *--------------------------------------------------------------------------*
```

Figure 4-10 Available Processors: Selecting a processor for the IOCP file

3. HCD displays a panel on which you enter information about the IOCP input data set to create (see Figure 4-11). Complete the following fields:
   - Title1: `PAVO2 IODF92`
   - IOCP input data set: `'ZNEXT01.IODF92.INPUT.IOCP'`
   - Input to Stand-alone IOCP: `Yes`
   - Job statement information: Complete this information for your installation.

```
*------------------------ Build IOCP Input Data Set -------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| IODF name . . . . . . . . . : 'ZNEXT01.IODF92.WORK'                         |
| Processor ID  . . . . . . . : PAV02                                         |
| Title1 . PAV02 IODF92                                                       |
| Title2 : ZNEXT01.IODF92.WORK - 2025-05-13 16:34                             |
|                                                                             |
| IOCP input data set                                                         |
| 'ZNEXT01.IODF92.INPUT.IOCP'______________________________                   |
| Input to Stand-alone IOCP?  Yes  (Yes or No)                                |
|                                                                             |
| Job statement information                                                   |
| //ZNEXT01H  JOB (SG248581),'FPINTO HCD',                                    |
| //          MSGCLASS=X,CLASS=A,MSGLEVEL=(1,1),                              |
| //          NOTIFY=&SYSUID                                                  |
| /*JOBPARM LINES=9999                                                        |
| //*                                                                         |
| //*                                                                         |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 4-11 Build IOCP Input Data Set: Data fields to be updated

4. Press Enter. HCD submits a batch job to create the data set.
5. Using an editor or browser tool of your choice, verify if the data set that you created exists and contains IOCP statements (see Example 4-3). In this example, we used Time Sharing Option (TSO). This data set is used as an input into the CMT.

Example 4-3 Top of the IOCP input dataset (truncated)

```
ID    MSG1='PAV02 IODF92',                                   *
      MSG2='ZNEXT01.IODF92.WORK - 2025-05-13 10:54',         *
      SYSTEM=(9175,1),LSYSTEM=PAV02,                         *
      TOK=('PAV02',008003321A0839311054393801251333F00000000,00*
      000000,'25-05-13','10:54:39','........','........')
RESOURCE PARTITION=((CSS(0),(PAV00A,A),(PAV00B,B),(PAV00C,C),(*
      PAV00D,D),(PAV00E,E),(PAV001,1),(PAV002,2),(PAV003,3),(P*
      AV004,4),(PAV005,5),(PAV006,6),(PAV007,7),(PAV008,8),(PA*
      V009,9),(*,F)),(CSS(1),(PAV01A,A),(PAV01B,B),(PAV01C,C),*
      (PAV01D,D),(PAV01E,E),(PAV01F,F),(PAV011,1),(PAV012,2),(*
      PAV013,3),(PAV014,4),(PAV015,5),(PAV016,6),(PAV017,7),(P*
      AV018,8),(PAV019,9)),(CSS(2),(PAV02A,A),(PAV02B,B),(PAV0*
      2C,C),(PAV02D,D),(PAV021,1),(PAV022,2),(PAV023,3),(PAV02*
      4,4),(PAV025,5),(PAV026,6),(PAV027,7),(PAV028,8),(PAV029*
      ,9),(*,E),(*,F)),(CSS(3),(PAV03A,A),(PAV03B,B),(PAV03C,C*
      ),(PAV03D,D),(PAV03E,E),(PAV03F,F),(PAV031,1),(PAV032,2)*
      ,(PAV033,3),(PAV034,4),(PAV035,5),(PAV036,6),(PAV037,7),*
      (PAV038,8),(PAV039,9)),(CSS(4),(PAV041,1),(PAV042,2),(PA*
      V043,3),(PAV044,4),(PAV045,5),(PAV046,6),(PAV047,7),(PAV*
```

Part of the TOK statement is now replaced with dots (see Example 4-4).

Example 4-4 IOCP file (TOK statement)

```
TOK=('PAV02',008003321A083931105439380125133F00000000,00*
      000000,'25-05-13','10:54:39','........','........')
```

These dots ensure that this IOCP file cannot be written to a processor and used for a Power on Reset (POR). This precaution is needed because this IOCP file was created from a validated work IODF and not a production IODF. IOCP files that can be used for a POR can be generated only from a production IODF.

> **Important:** When an IOCP file is exported by using HCD from a validated work IODF, it must be imported back into HCD after the *channel IDs* (CHIDs) are complete by using the CMT. The IOCP file cannot be used directly by IOCP until the CHIDs are added.

6. Download this IOCP file from z/OS to the CMT workstation. Use a workstation file transfer facility such as SFTP or the one in the IBM Personal Communications Workstation Program, or any equivalent 3270 emulation program. Be sure to use `TEXT` as the transfer type. In this example, the file is named `ZNEXT01.IODF92.INPUT.IOCP`.

4.3 Assigning CHIDs to CHPIDs by using the CMT

In this section, you use the IOCP statements from HCD and the 9175 order process file (CFReport). Use the CMT to assign CHIDs to each of the CHPIDs for the 9175.

For this process, the CMT must be downloaded. For more information about downloading and installing the CMT, see 2.2.3, “IBM Z Connectivity Mapping Tool” on page 12. If the CMT is already installed, verify if the latest updates are installed.

The version of the CMT supporting IBM z17 that is used for the following captures is Version 1.0.0. Check for the latest versions by going to IBM Resource Link.

For more information, see the *CMT User’s Guide,* GC28-7058 also available on IBM Resource Link.

Use the CMT to complete the following steps:

1. Import the CFReport file into the CMT.
2. Import the IOCP file into the CMT.
3. Resolve CHPIDs with a CHID conflict.
4. Process the hardware resolution.
5. Manually resolve the coupling link CHPIDs.
6. Set the priority for single-path control units (CUs) and other CUs that override the CMT default priorities and Automatic Mapping.
7. Resolve the CHPIDs that are not connected to CUs.
8. Create the CMT reports.
9. Create an updated IOCP file for transfer back into the IODF file.

4.4 Importing the CFR file and the IOCP input file into CMT

To import the CFReport file into the CMT, complete the following steps:

1. Start the CMT on your workstation.

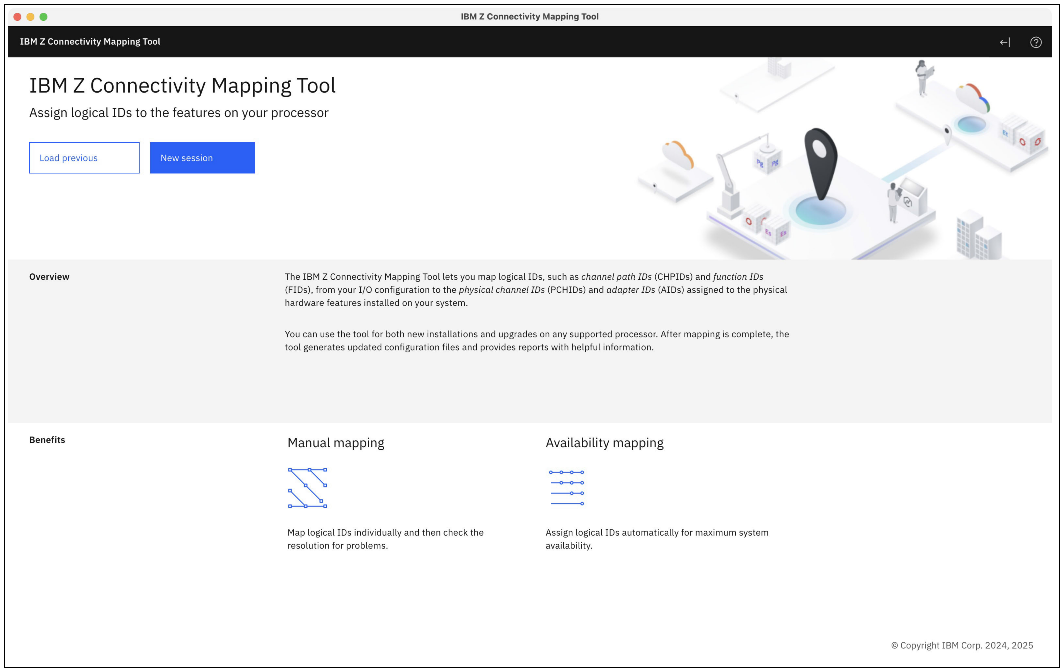

Figure 4-12 Shows the CMT starting page

2. The CMT asks for a session name, the CFR report and the IOCP input file. In our example, we used `PAV022z17` as the session name (Figure 4-13).

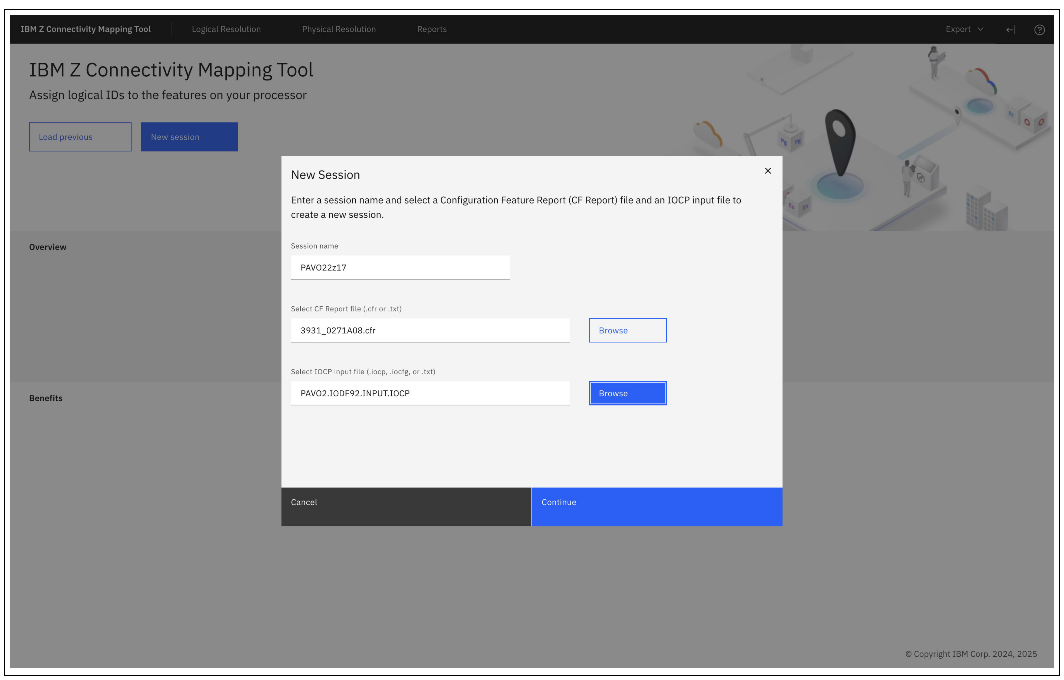

Figure 4-13 Creating a CMT Session and loading the files

To import the validated 9175 IOCP file into the CMT, enter the values into the create session window, see Figure 4-13, click on `CONTINUE`.

> **Note:** The functionality to produce location reports based on only the CFR file will be made available in the next version of the IBM Z CMT.

4.5 Building the initial Overview

The `NEW SESSION` window appears, containing the information about the new session and the input files, see Figure 4-14 for details. Since this is an upgrade scenario, CMT detects both processor types and asks, which representation should be used. We selected `Proposed configuration` so the CHIDs for the new processor type can be assigned by CMT.

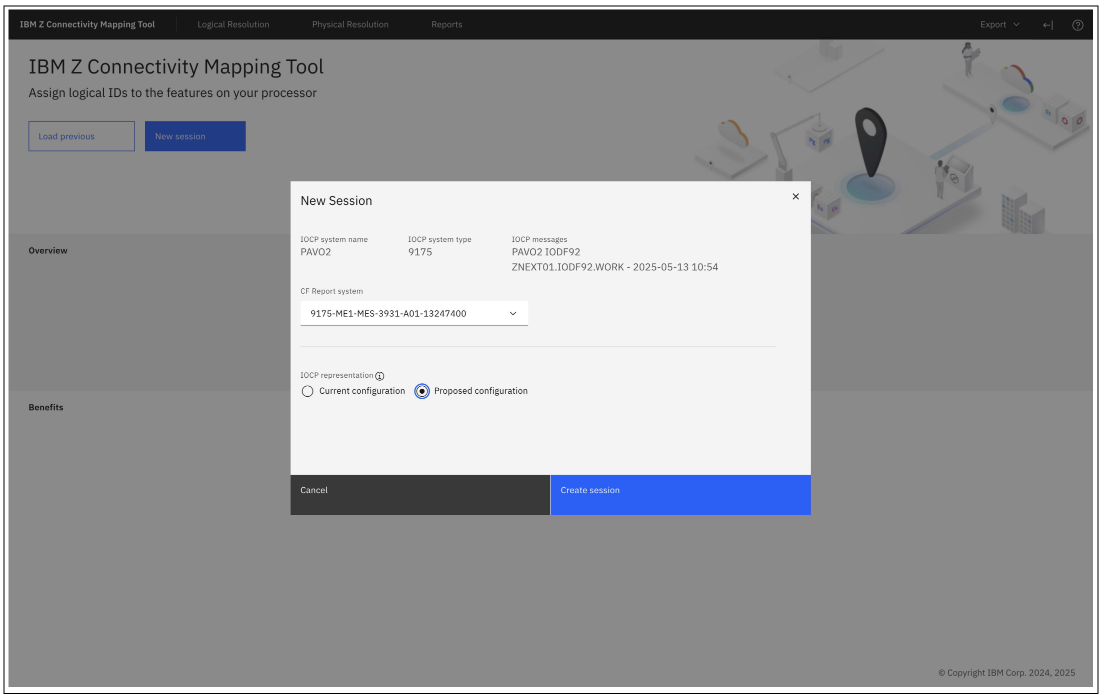

Figure 4-14 New Session window with session information

Click on the `Create session` button. The Logical resolution window appears. See Figure 4-15 on page 76.

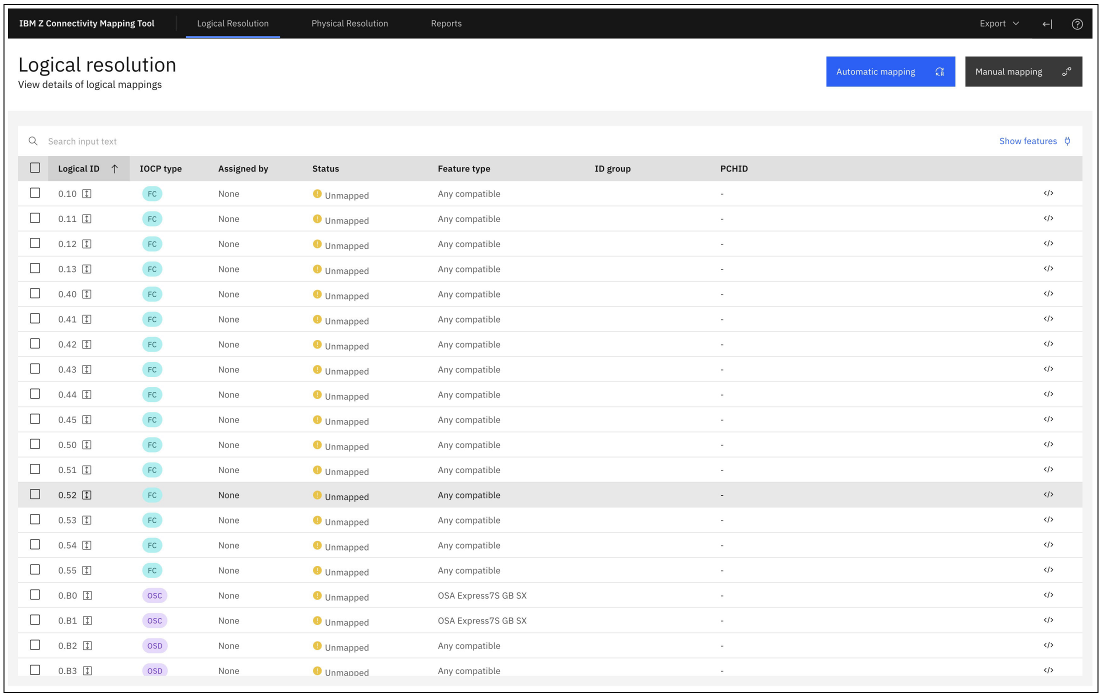

Figure 4-15 Logical resolution window

The CMT shows the information from the CFReport file and the IOCP file in the `Logical resolution` panel by default, the top of the panel you can select one of three main panels (see Figure 4-15):

- Logical resolution
- Physical resolution
- Reports

On the `Logical resolution` panel the `Logical IDs` show the logical channel subsystem and the channel path in the format (CSS.CHPID) they are displayed as well as the `IOCP type` (channel type derived from the IOCP input file) and the assigned feature type if possible.

The `Physical resolution` panel shows a table with the physical location of the ordered features and, if possible, assigned Logical IDs to PCHIDs. See Figure 4-16 on page 77 for details.

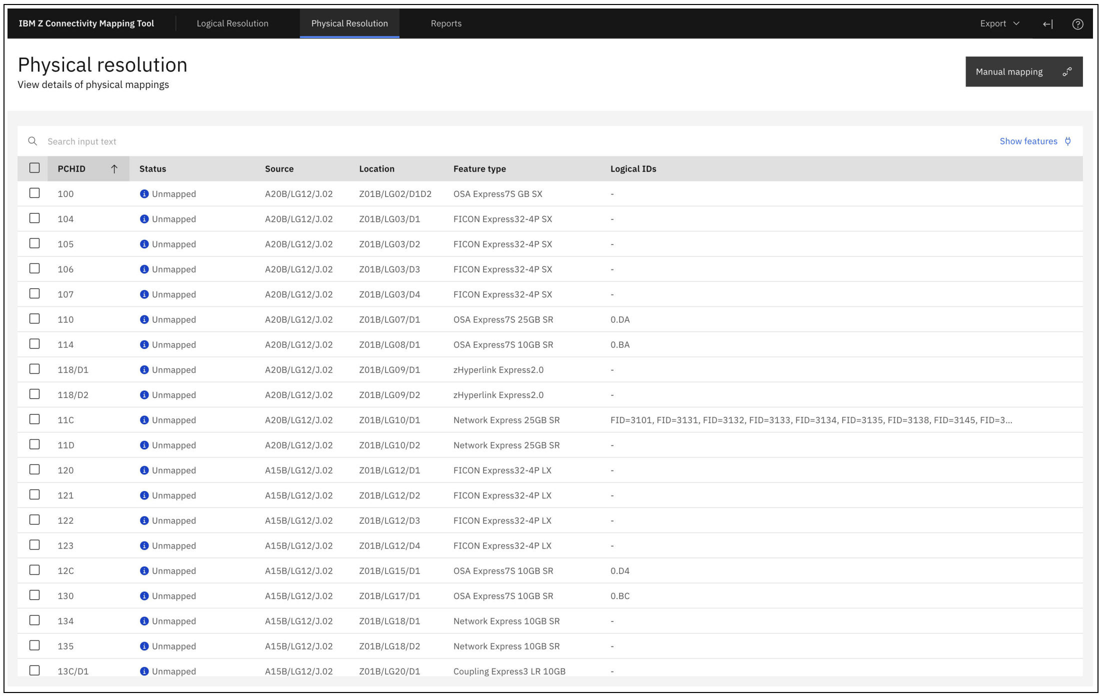

Figure 4-16 Physical resolution showing the availability of physical resources

On the upper right corner the `Show features` button summarizes the number of used and available features, see Figure 4-17.

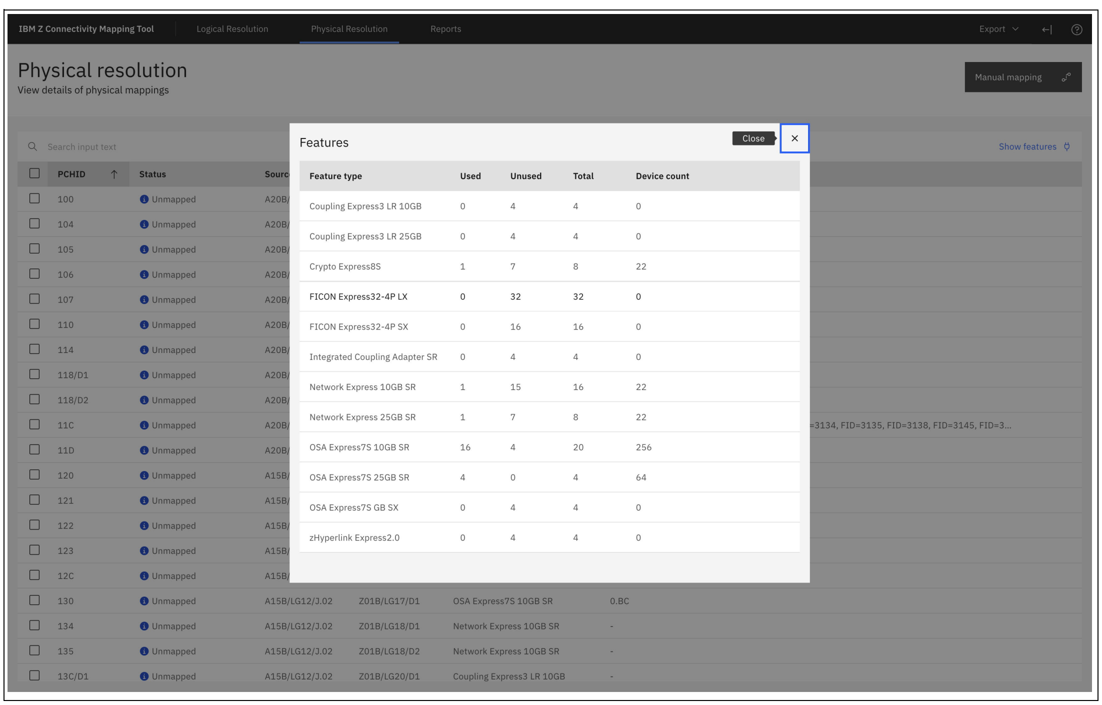

Figure 4-17 Show features window

.

4.6 Assigning missing feature types

In the `Logical Resolution` panel, the CMT shows the auto detected features as well as all the channel types and the `Logical IDs` (CHPIDs) as derived from the IOCP. In the next step we will assign the appropriate features.

We select all the `Logical IDs` with `IOCP type FC` and assign the FICON Express32 - 4P LX adapter card. First select all Logical ID, then on the top right corner click on `Assign feature type`, see Figure 4-18.

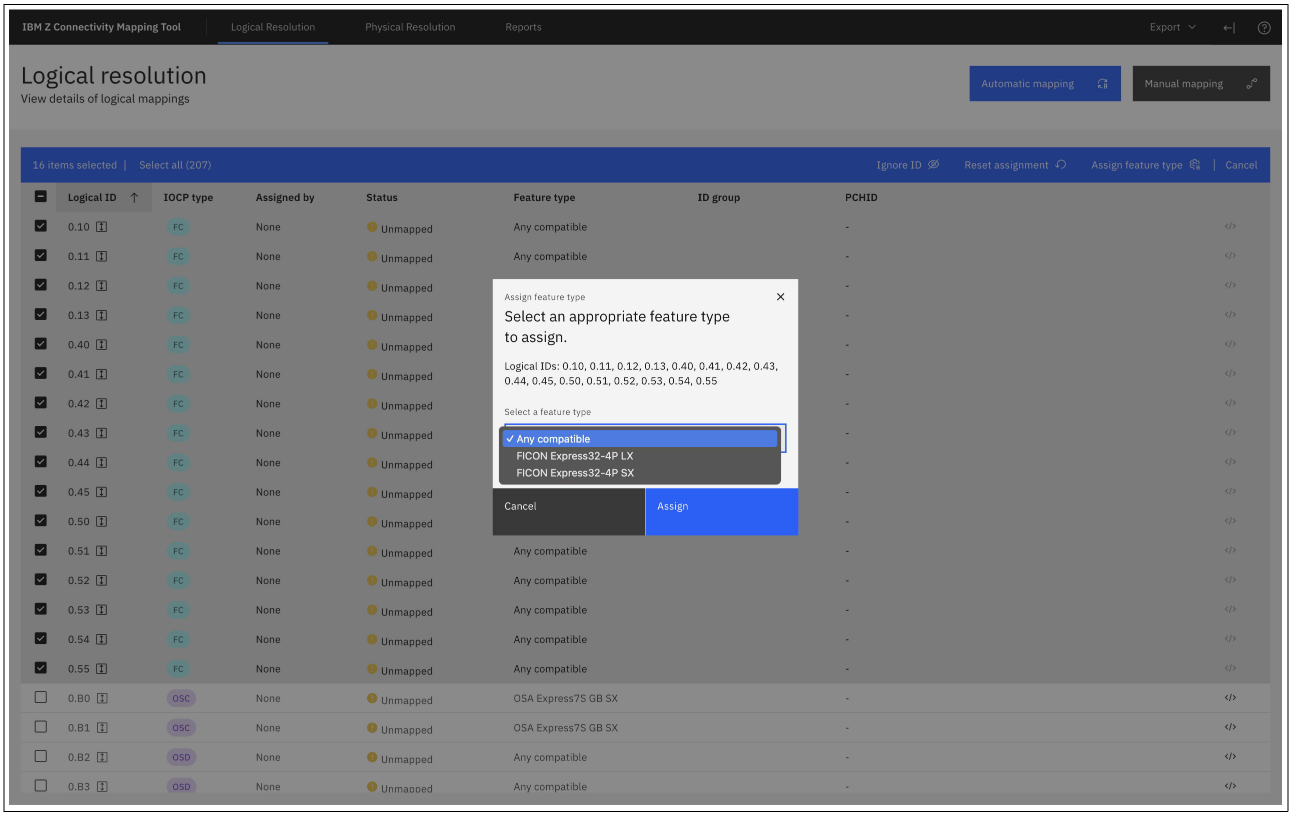

Figure 4-18 Assign features to selected `IOCP types` and `Logical IDs`

CMT assigns the selected feature type to the Logical ID. See Figure 4-19 on page 79.

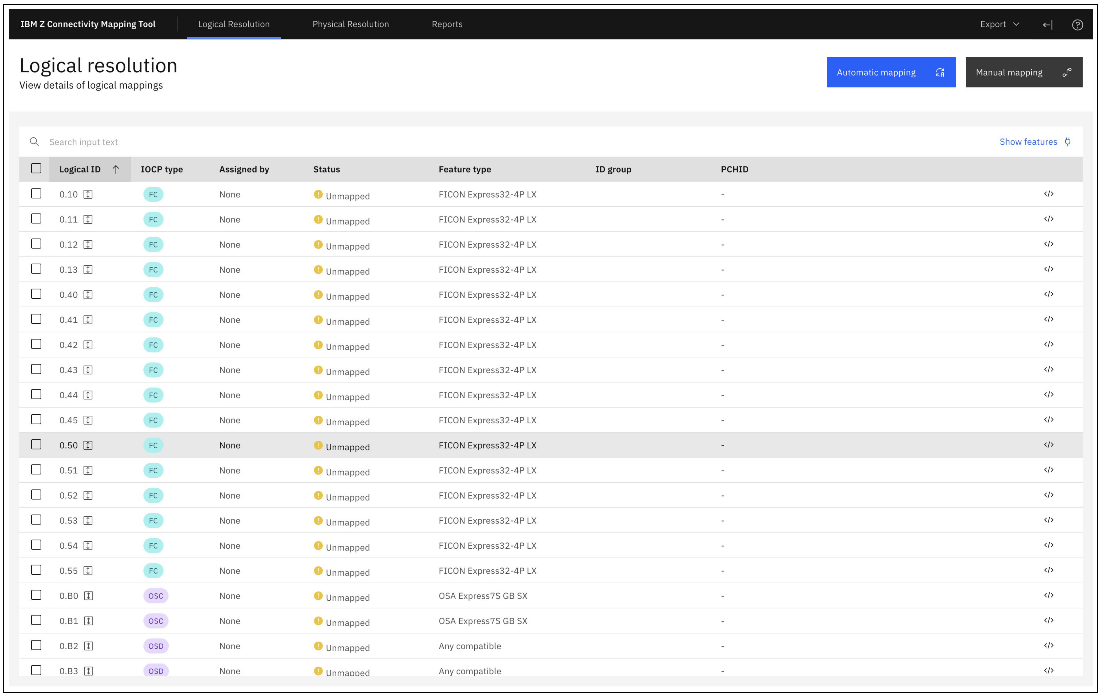

Figure 4-19 Logical Resolution panel with selected features assigned

Because the 9175 configuration contains both Network Express 10Gb SR/LR and Network Express 25GB SR/LR adapter cards, we have to assign the appropriate adapter type to the IOCP type manually. Use the same process and first select the Logical IDs you plan to use the same feature for. Then click on `Assign feature type` on the upper right corner of the panel to get a list of available features for the IOCP type. See Figure 4-20.

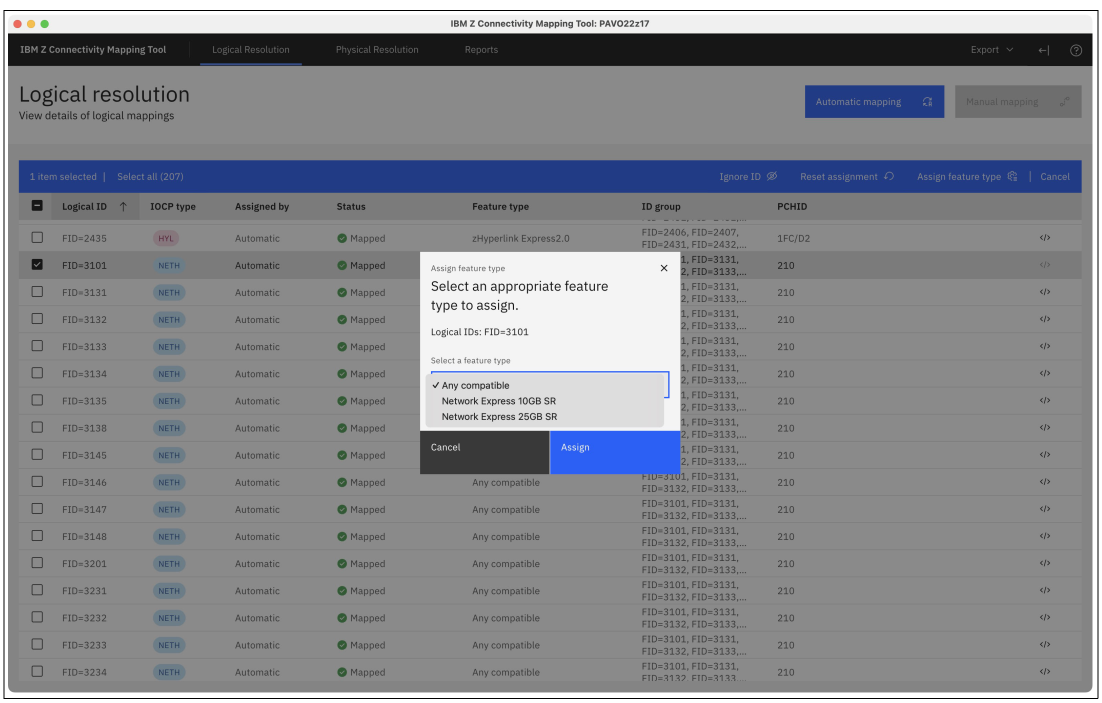

Figure 4-20 Assign appropriate feature type to `IOCP type`

4.7 Resetting Incompatible (Hardware - I/O) Entries

In this example, CMT is used to manually assign a physical AID/Port designation to CL5 coupling links. In Figure 4-21 an example of a miss configured coupling links is shown.

The AID/Port designation has been derived from the IBM z16 configuration not matching the new IBM z17 configuration.

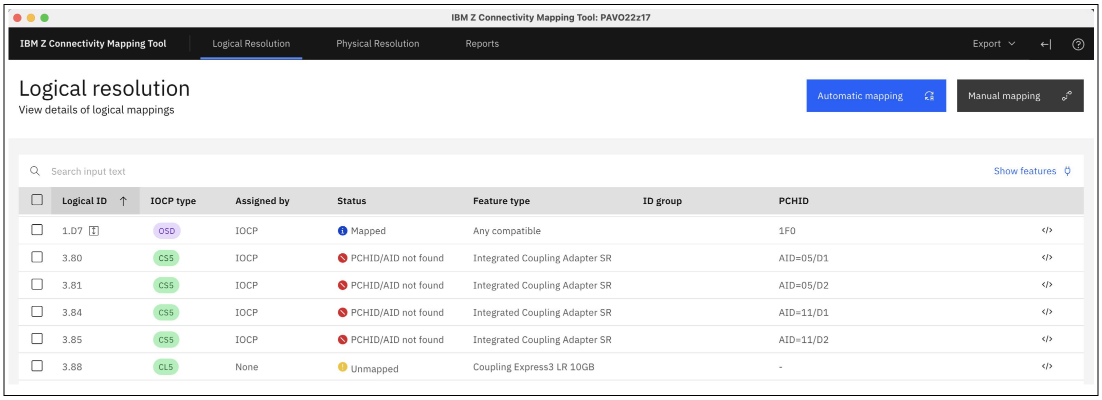

Figure 4-21 CS5 Coupling Links with definition errors

The first step is to remove the current assignment. This is done by selecting the `Logical ID` on the left of the panel and then selecting `Reset assignment` on the upper right corner as shown in Figure 4-22 on page 80.

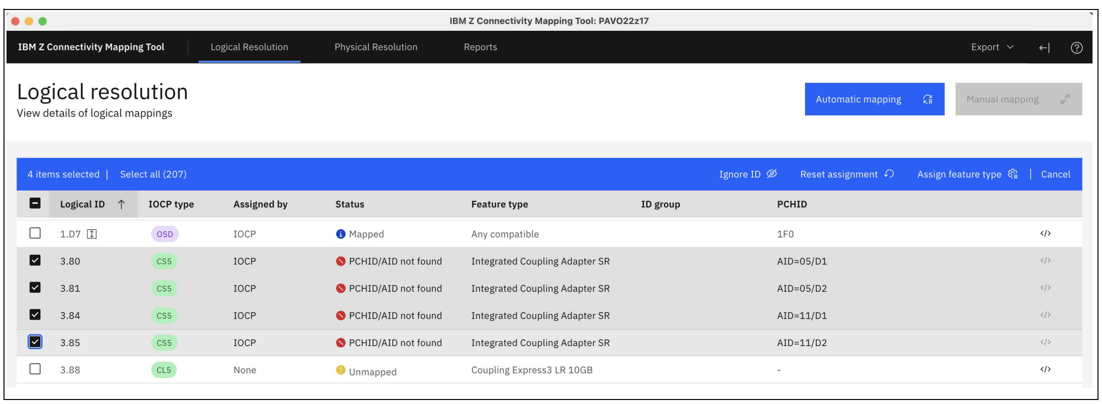

Figure 4-22 Preparing to reset assignments

After resetting the assignments re-select the Logical ID and click on Manual Mapping the resulting panel let’s you assign the AID/Port combination present on the upgraded 9175 configuration, see Figure 4-23.

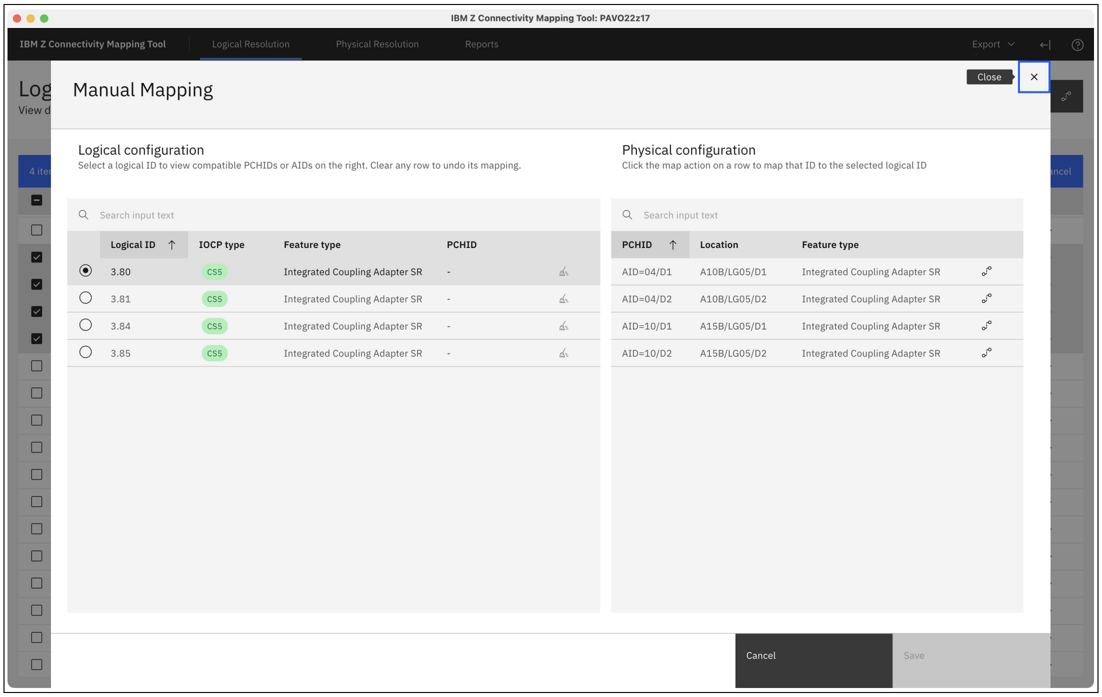

Figure 4-23 Assign AID/Port to coupling links

Assign the AID/Port to the appropriate `Logical ID` by clicking on the `MAP` symbol on the right of the `Manual Mapping` panel. Then click on `Save` on the lower right side of the panel. The correct AID/Port combination have now been assigned, see Figure 4-24 on page 81.

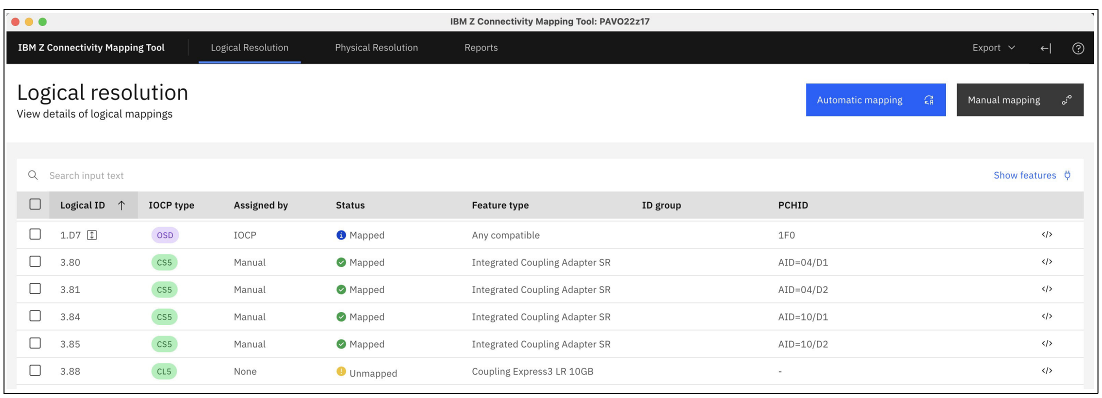

Figure 4-24 Successful AID/Port assignment

> **Note:** For more information about these error messages, see the *Connectivity Mapping Tool User's Guide, GC28-7058*.

4.7.1 Assigning PCHID entries to FUNCTION IDs (FID)

`FIDs` have been ported from the previous configuration and need to be re-assigned to new `PCHIDs` because of new HyperLink Adapter cards. See Figure 4-25.

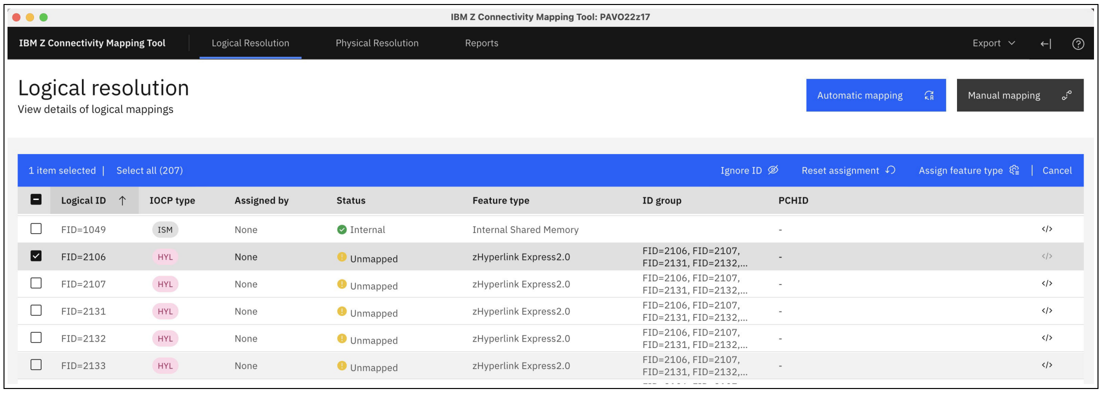

Figure 4-25 HyperLink Functions missing `PCHIDs`

Select the FID with the lowest number, and click on `Manual mapping` on the upper right corner of the panel, then select the `PCHID` for this `FUNCTION` and click `Save`. CMT will automatically assign the `PCHID` for the whole group of `FIDs`, see Figure 4-26 on page 82.

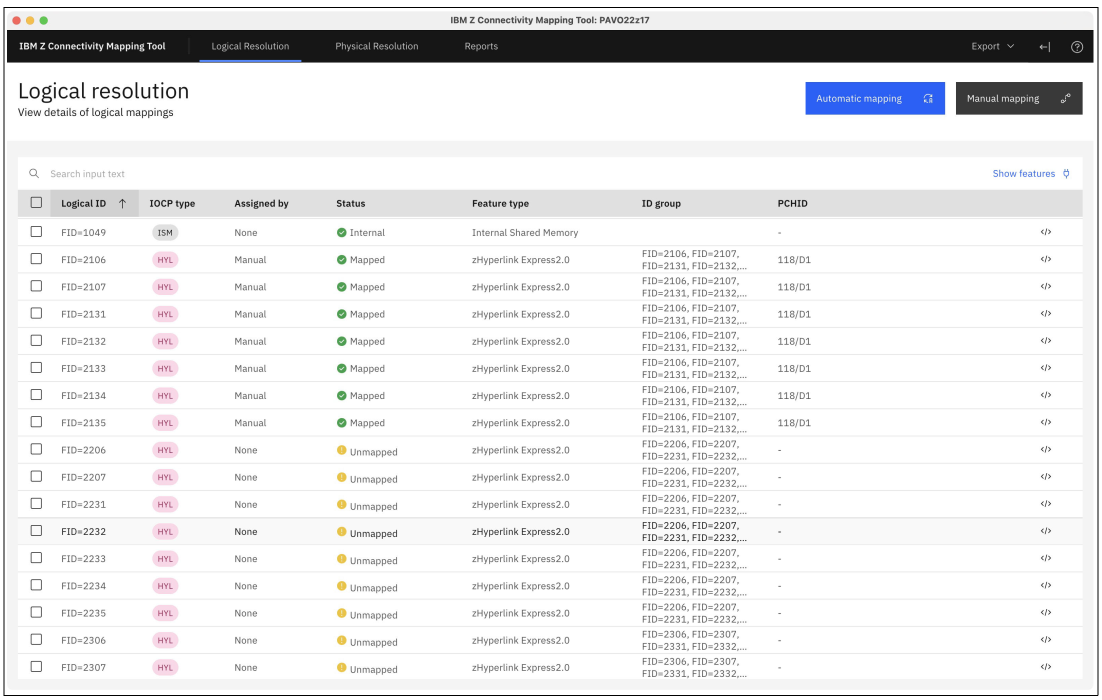

Figure 4-26 `PCHID` assigned to a whole group of `HYL FIDs`.

4.8 Creating CMT reports

The CMT offers built-in reports, which are available from the top of the panel clicking `Reports`. You can also export the information as a CSV file from the report by clicking `Download CSV` on the lower right corner of each tile. Figure 4-27 shows the options to create a Preview Report or download a report.

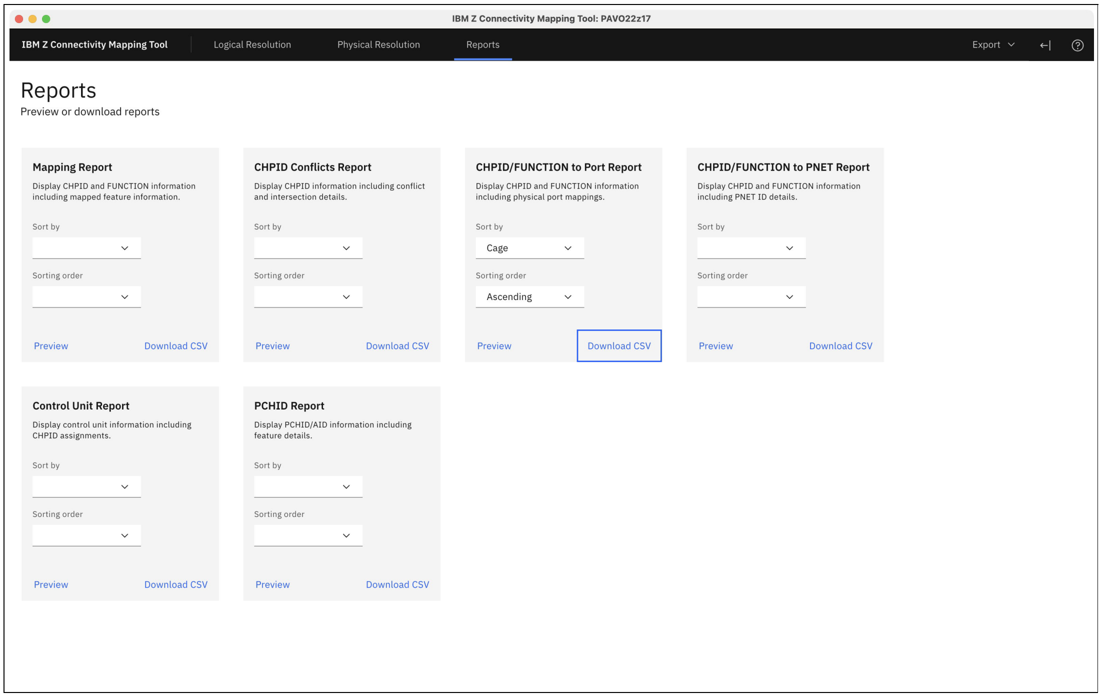

Figure 4-27 Preview or download reports panel

Click on the `Sort by` and `Sort order` buttons in the appropriate tile. Then either click on `Preview Report` or `Download CSV` to display the results or export them as a CSV file. See Figure 4-28 for selection criteria for a CHPID report.

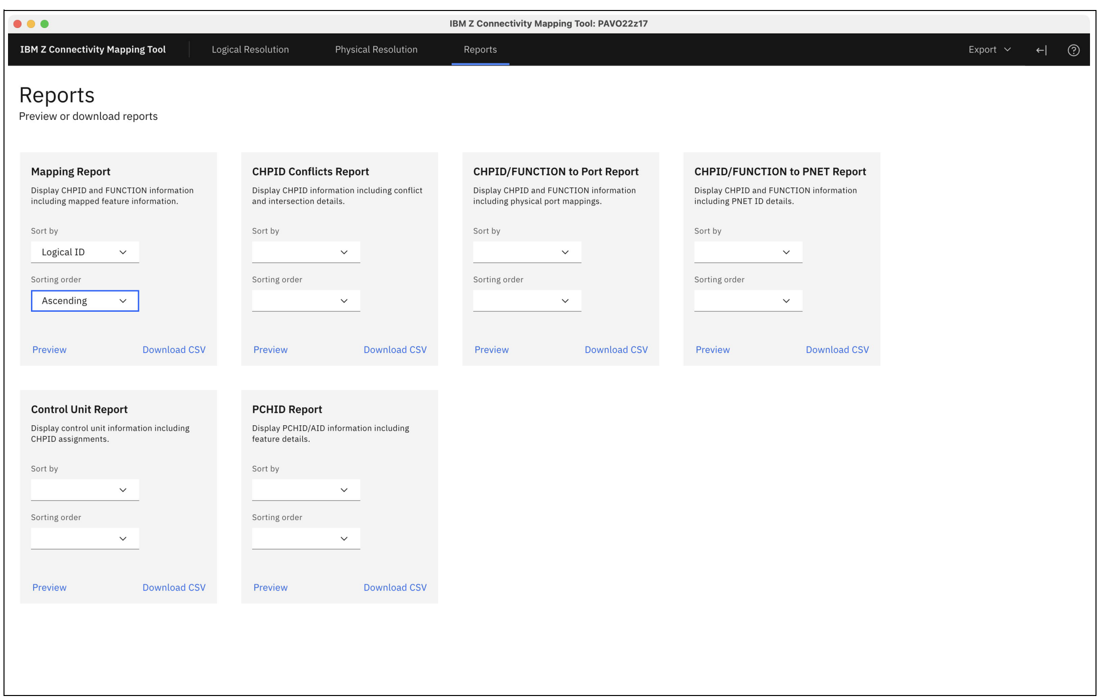

Figure 4-28 Selection of `Mapping Report` sorted by `Logical ID`

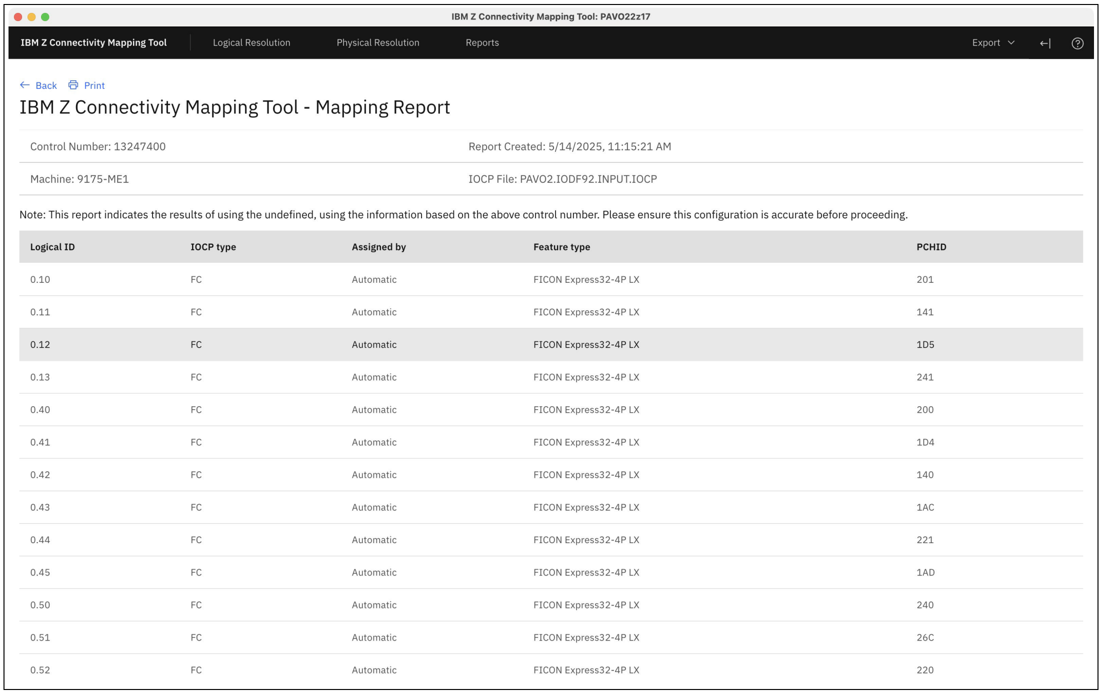

Figure 4-29 Preview of CHPID mapping report

If requested, the report can also be printed or exported to PDF. Click on the `Print` symbol on the top left corner of the panel.

For simplicity, only three reports are described in this example:

- The CHPID Report
- The Port Report sorted by location
- The CHPID to Control Unit Report

However, all built-in reports are created the same way.

4.8.1 CHPID Report

To create the CHPID Report, complete the following steps:

1. On the CMT main panel select `Reports` from the top.
2. In the `Mapping Report` tile, on the `Sort by` button select `Logical ID` and on the `Sort by` button select `Ascending`. See Figure 4-28 on page 83 and Figure 4-29 for the selection process and the resulting report.

CHPID to Port Report sorted by location

To create the Port Report sorted by location, on the `Reports` panel select the CHPID/FUNCTION to Port Report tile and select Cage on the Sort by button and Ascending on the Sort by button. See Figure 4-30 on page 85.

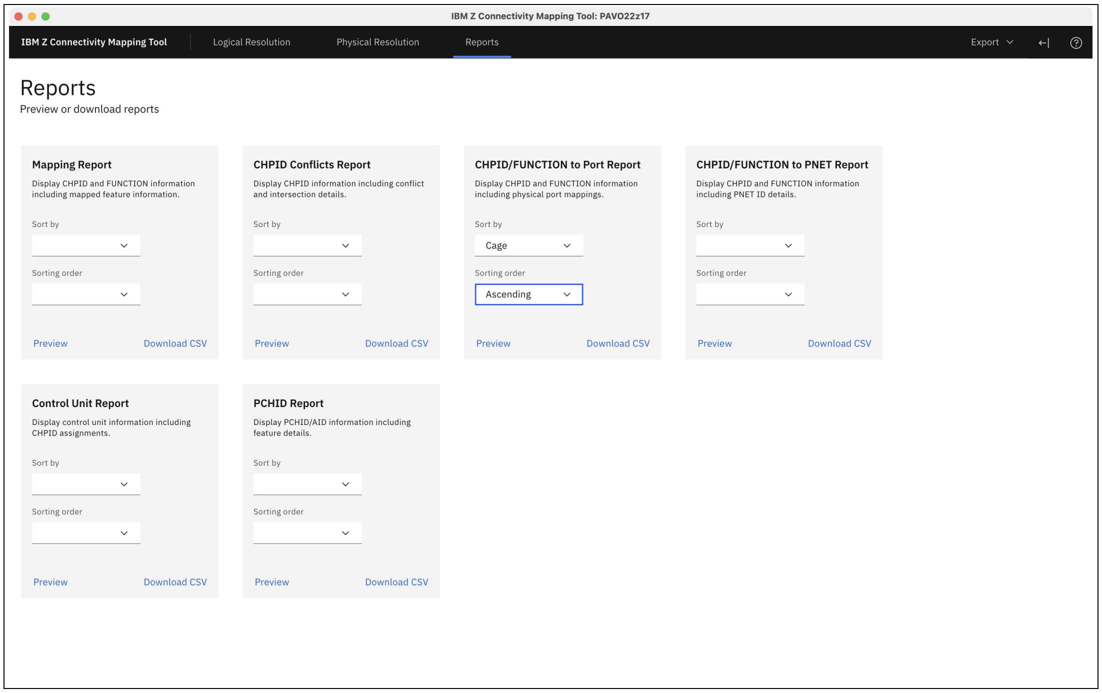

Figure 4-30 Selection for CHPID report sorted by physical location

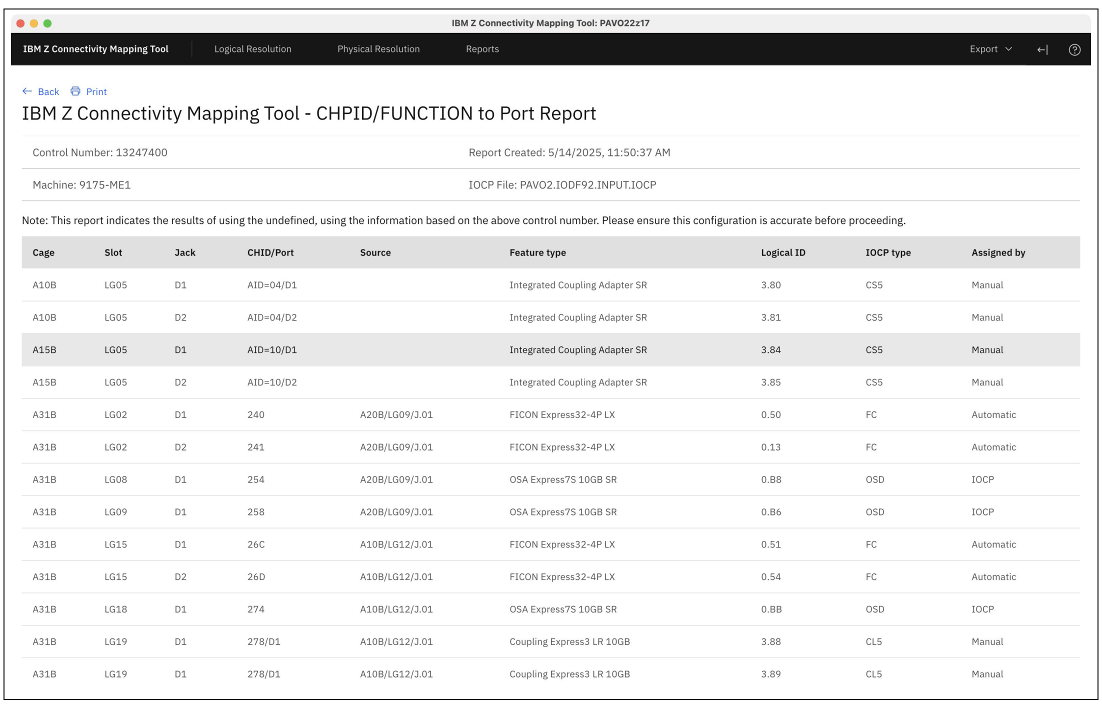

Figure 4-31 Sample output CHPID mapping report sorted by physical location

The person who installs the I/O cables during system installation needs one of these reports. The Port Report sorted by location is preferable. The installer can use this report to help with labeling the cables. The labels must include the CHID or cage, slot, and port information before system delivery. See Figure 4-31 for details.

CHPID to CU Report

The Control Unit report provides an overview of all the installed Control Units, their types and the attached CHPIDs as well as the CSS source. To generate the report, from the Reports panel select the CSS option on the Sort by button and the Ascending option in the Sorting order button in the Control Unit Report tile. See Figure 4-32.

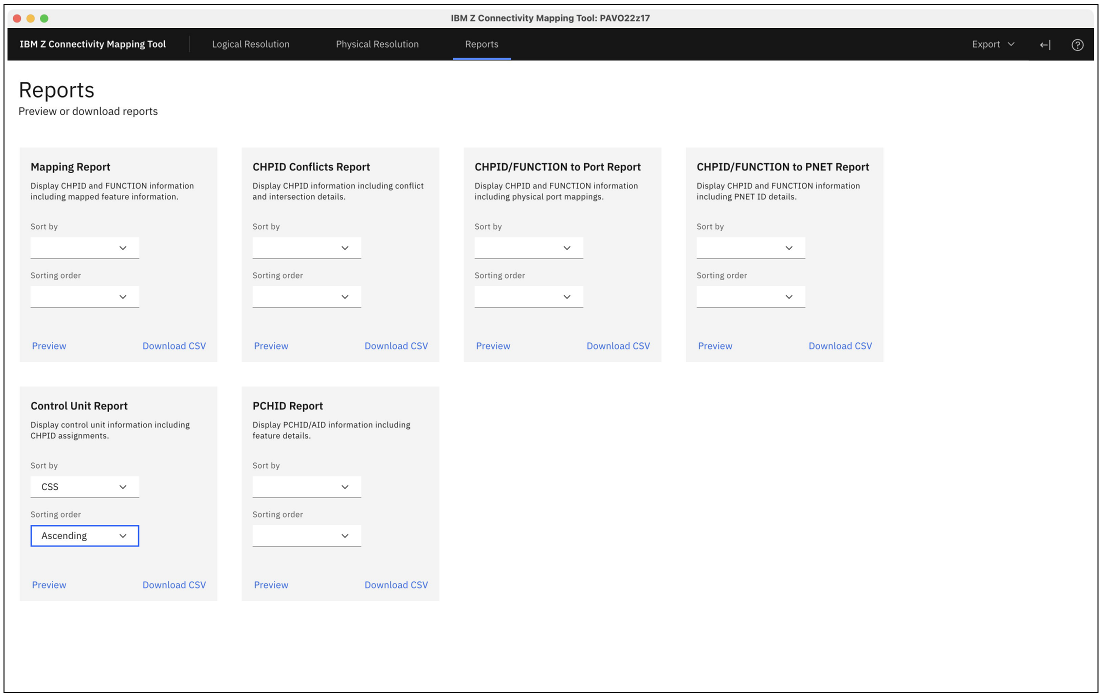

Figure 4-32 Selection options for the Control Unit Report

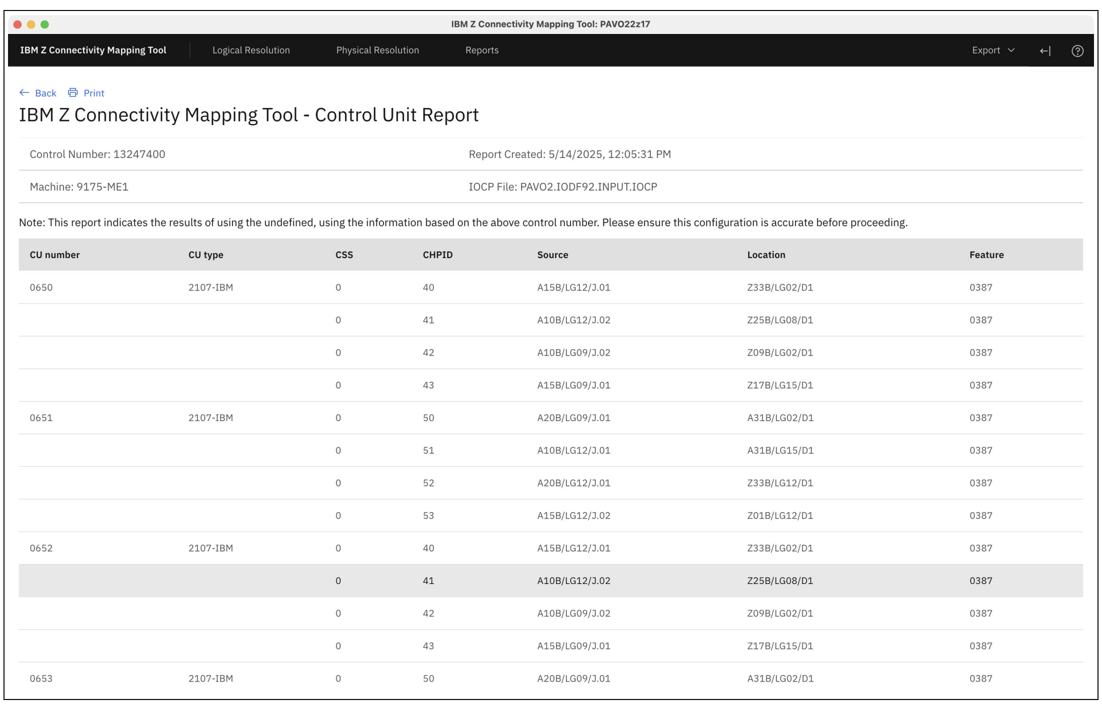

Figure 4-33 Example of a Control Unit Report

4.9 Creating an updated IOCP file

After all the indicated conditions have been resolved, CMT can be used to create an updated IOCP file including the inserted PCHIDs. The file can then be uploaded into the LPAR where HCD is running. Once uploaded, the file can be used to be merged with the IODF used to produce the IOCP input file for CMT.

To create the updated IOCP, complete the following steps:

1. From the main panel click on Export on the top right corner of the panel. See Figure 4-34.

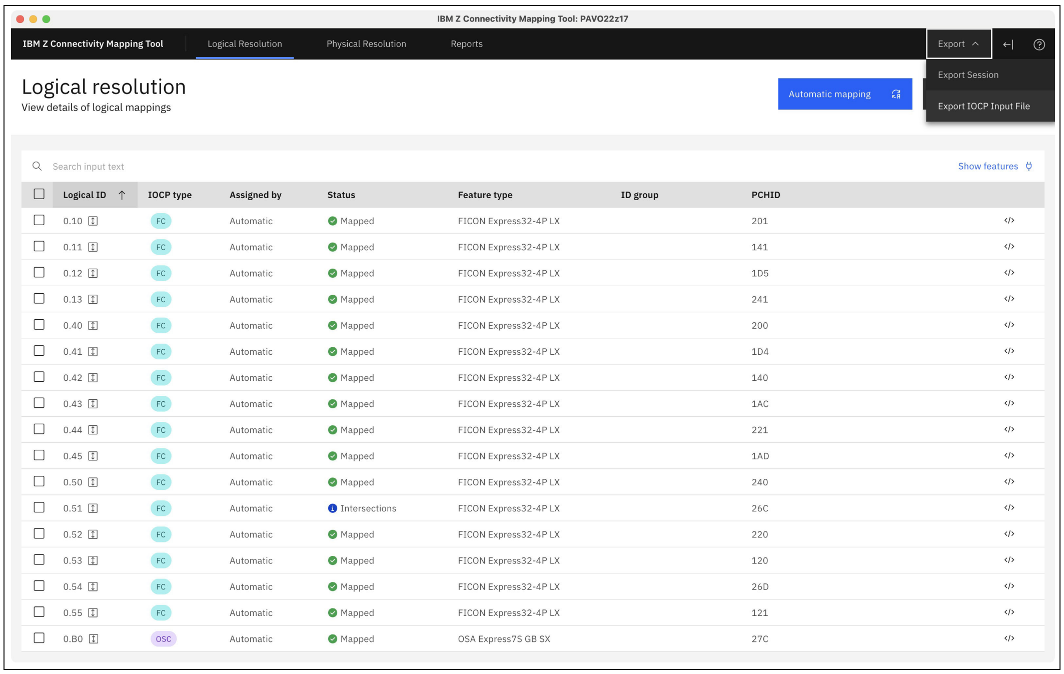

Figure 4-34 Prepare to export the current session or the updated IOCP file

Once Export IOCP Input File has been selected, on the next panel the destination can be selected. See Figure 4-35 on page 88

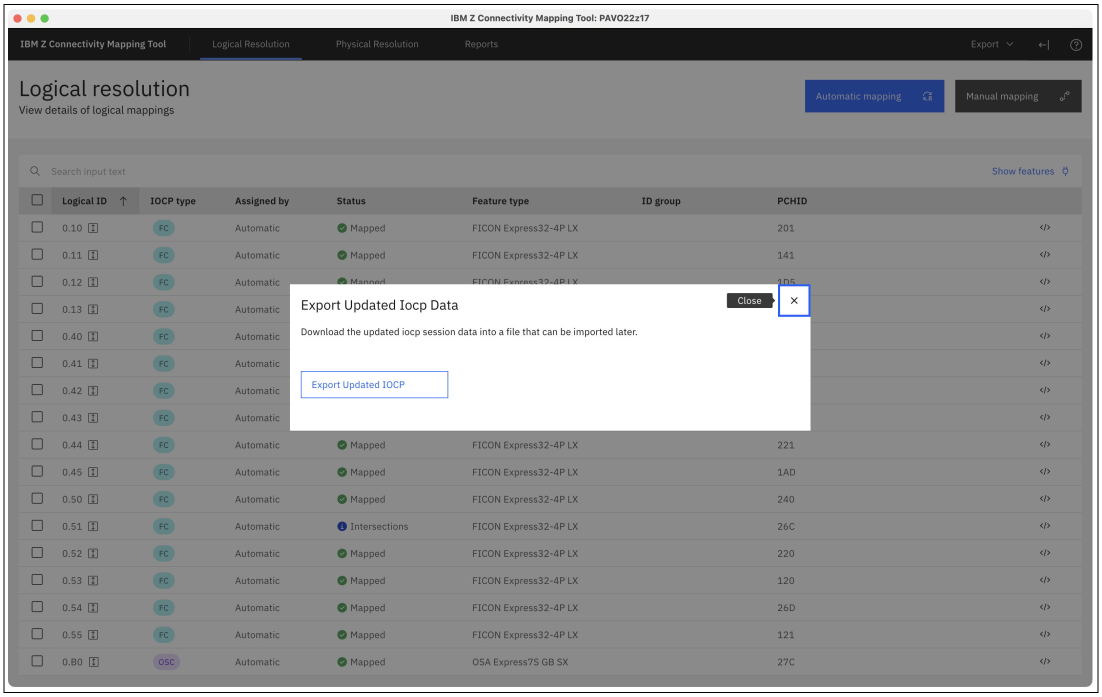

Figure 4-35 Execute export and define name and location of the file

Click on the Export Updated IOCP button and enter the Export Path and IOCP Name for the IOCP output file (see Figure 4-35).

> **Requirement:** This file must be uploaded to the z/OS image on which you have the work IODF that you used previously to create the IOCP input data set.

Saving the session for later reference can be very useful. On the Logical resolution panel, click Export on the top right corner of the panel, click on Export Session. On the Export CMT Session panel click on Export Session, enter the Export Path and Session Name. See Figure 4-34 on page 87.

4.10 Additional steps and processes

You might want to perform a PCHID migration before building a production IODF. For more information, see *Connectivity Mapping Tool User's Guide,* GC28-7024.

For your next steps, go to Chapter 5, “Building the production input/output definition file and setting up the central processor complex” on page 89.

5

Building the production input/output definition file and setting up the central processor complex

This chapter describes the tasks that are needed to build a production input/output definition file (IODF) and set up the central processor complex (CPC).

Naming: The IBM z17 systems that are targeted by this publication consist of IBM z17 ME1. Throughout this chapter, we might refer to these machines as IBM z17.

Note: Although IBM z15 T01 (8561) and IBM z16 A01 (3931) systems are upgradable to IBM z17 ME1 (9175), this chapter describes the configuration steps to upgrade an existing IBM z16 to the IBM z17.

This chapter includes the following topics:

- Building the new production IODF
- Writing the input/output configuration program to the old CPC by using HCD
- Creating a reset profile on the Support Element
- Creating an image profile on the Support Element
- Performing a Power on Reset on the new CPC
- Building and verifying load (IPL) profiles
- Building and verifying LOADxx members in SYS#.IPLPARM
- Communicating information about the new CPC

5.1 Building the new production IODF

To use the definitions that were updated in Hardware Configuration Definition (HCD), create a production IODF from your work IODF. Then, write the IODF to the input/output configuration data set (IOCDS) by using Write IOCDS in preparation for the upgrade.

Complete the following steps:

1. From the HCD main panel, select option `2. Activate or process configuration data`. Confirm the `I/O definition file` name at the bottom of the panel. (see Figure 5-1).

```
z/OS V3.1 HCD
Command ===> ________________________________________________________________

                           Hardware Configuration

Select one of the following.

2   0.  Edit profile options and policies
    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 . . . 'ZNEXT01.IODF89.WORK'                +
```

Figure 5-1 Hardware Configuration: Activate or process configuration data

2. The Activate or Process Configuration Data panel opens (see Figure 5-2). Select option `1. Build production I/O definition file`, and then press Enter.

```
+------ Activate or Process Configuration Data ------+
|                                                     |
|                                                     |
| Select one of the following tasks.                  |
|                                                     |
| 1_  1.  Build production I/O definition file        |
|     2.  Build IOCDS                                 |
|     3.  Build IOCP input data set                   |
|     4.  Create JES3 initialization stream data      |
|     5.  View active configuration                   |
|     6.  Activate or verify configuration            |
|         dynamically                                 |
|     7.  Activate configuration sysplex-wide         |
|     8.  (no longer supported)                       |
|     9.  (no longer supported)                       |
|     10. Build I/O configuration data                |
|     11. Build and manage processor cluster          |
|         IOCDSs, IPL attributes and dynamic I/O      |
|         changes                                     |
|     12. Build validated work I/O definition file    |
|                                                     |
|  F1=Help     F2=Split    F3=Exit     F9=Swap        |
| F12=Cancel                                          |
```

Figure 5-2 Activate or Process Configuration Data: Build production I/O definition file

3. HCD displays the Message List panel (see Figure 5-3). Verify that you have only severity "W" (warning) messages and that they are normal for the configuration. Correct any other messages that should not occur and try to build the production IODF again. Continue this process until you have no messages that indicate problems.

```
+------------------------------- Message List --------------------------------
|   Save  Query  Help
| --------------------------------------------------------------------------
|                                                                Row 1 of 99
| Command ===> ___________________________________________ Scroll ===> CSR
|
| Messages are sorted by severity. Select one or more, then press Enter.
|
| / Sev Msg. ID  Message Text
| _ W   CBDG098I For operating system DBSV4SU4 and device type OSA the
| #              default of LOCANY=YES is not used for following device
| #              groups: 1910,16 1930,16
| _ W   CBDG098I For operating system DBSV5SU4 and device type OSA the
| #              default of LOCANY=YES is not used for following device
| #              groups: 1910,16 1930,16
| _ W   CBDG098I For operating system DBSV6SU4 and device type OSA the
| #              default of LOCANY=YES is not used for following device
| #              groups: 1910,16 1930,16
| _ W   CBDG098I For operating system PERF4SU4 and device type OSA the
| #              default of LOCANY=YES is not used for following device
|  F1=Help        F2=Split       F3=Exit        F4=Prompt      F5=Reset
|  F7=Backward    F8=Forward     F9=Swap       F10=Actions    F12=Cancel
| F13=Instruct   F22=Command
+-----------------------------------------------------------------------------
```

Figure 5-3 Message List: Building a production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (see Figure 5-4). Complete the Production IODF name and Volume serial number fields, and then press Enter.

```
+-------------- Build Production I/O Definition File ---------------+
|                                                                    |
|                                                                    |
| Specify the following values, and choose how to continue.          |
|                                                                    |
| Work IODF name . . . : 'ZNEXT01.IODF89.WORK'                       |
|                                                                    |
| Production IODF name . 'SYS9.IODF81'________________________       |
| Volume serial number . IODFPK  +                                   |
|                                                                    |
| Continue using as current IODF:                                    |
| 2   1.  The work IODF in use at present                             |
|     2.  The new production IODF specified above                    |
|                                                                    |
|                                                                    |
|                                                                    |
|                                                                    |
+--------------------------------------------------------------------+
```

Figure 5-4 Build Production I/O Definition File: Data fields to be updated

6. The Define Descriptor Fields panel opens (see Figure 5-5). Press Enter to accept the descriptor fields that are selected by HCD, or enter different values and then press Enter.

```
+-------------------- Define Descriptor Fields --------------------+
|                                                                  |
|                                                                  |
| Specify or revise the following values.                          |
|                                                                  |
| Production IODF name  . : 'SYS9.IODF81'                          |
|                                                                  |
| Descriptor field 1  . . . SYS9                                   |
| Descriptor field 2  . . . IODF81                                 |
|  F1=Help      F2=Split     F3=Exit      F5=Reset     F9=Swap     |
| F12=Cancel                                                       |
+------------------------------------------------------------------+
```

Figure 5-5 Define Descriptor Fields: Data fields to be updated

HCD displays the following message, which indicates that the production IODF was successfully created:

```
Production IODF SYS9.IODF81 created.
```

To implement the configuration on the 3931 processor in preparation for its upgrade to a 9175 processor, go to 5.2, “Writing the input/output configuration program to the old CPC by using HCD” on page 93.

5.2 Writing the input/output configuration program to the old CPC by using HCD

Now that you have a production IODF that is named `SYS9.IODF81`, you can now write the input/output configuration program (IOCP) data from the IODF to the IOCDS on the CPC that you want to upgrade (for example, PAVO).

The IOCDS are available for Power on Reset (POR) after the processor is upgraded.

To update the IOCDS by using HCD option 2.11, complete the following steps:

1. From the HCD main panel, select option `2. Activate or process configuration data` (see Figure 5-6). Ensure that the IODF is the production IODF that was created in 5.1, “Building the new production IODF” on page 90, and then press Enter.

```
 z/OS V3.1 HCD
Command ===> ________________________________________________________________

                           Hardware Configuration

Select one of the following.

2   0.  Edit profile options and policies
    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 . . . 'SYS9.IODF81'                          +
```

Figure 5-6 Hardware Configuration: Activate or process configuration data

2. The Activate or Process Configuration Data panel opens (see Figure 5-7). Select option `11. Build and manage processor cluster IOCDSs, IPL attributes and dynamic I/O changes`, and then press Enter.

```
+------ Activate or Process Configuration Data ------+
|                                                     |
|                                                     |
| Select one of the following tasks.                  |
|                                                     |
| 11  1.  Build production I/O definition file        |
|     2.  Build IOCDS                                 |
|     3.  Build IOCP input data set                   |
|     4.  Create JES3 initialization stream data      |
|     5.  View active configuration                   |
|     6.  Activate or verify configuration            |
|         dynamically                                 |
|     7.  Activate configuration sysplex-wide         |
|     8.  (no longer supported)                       |
|     9.  (no longer supported)                       |
|     10. Build I/O configuration data                |
|     11. Build and manage processor cluster          |
|         IOCDSs, IPL attributes and dynamic I/O      |
|         changes                                     |
|     12. Build validated work I/O definition file    |
|                                                     |
|  F1=Help     F2=Split    F3=Exit     F9=Swap        |
| F12=Cancel                                          |
+-----------------------------------------------------+
```

Figure 5-7 Activate or Process Configuration data: Build and manage processor cluster IOCDSs, IPL attributes, and dynamic I/O changes

This example assumes that you have connectivity to the 3931 processor that is being upgraded over the Hardware Management Console (HMC) local area network (LAN) to write an IOCDS.

If the CPC being upgraded is not accessible from the HMC LAN, create a IOCP file from HCD, then use the stand-alone IOCP process to update the IOCDS.

You can create an IOCP file by using the same process that you used to create an IOCP file for the IBM Z Connectivity Mapping Tool.

Tip: The Support Element (SE) can read an IOCP file that is written to a USB flash memory drive.

3. The Processor Cluster List panel opens (see Figure 5-8). In the list, select the 3931 processor being upgraded by typing a forward slash (/) to update one of its IOCDSs, and then press Enter.

```
 Processor Cluster List                 Row 1 of 6
Command ===> _______________________________________________ Scroll ===> PAGE

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# IBM390PS.HYDRA     8562   LT2
# IBM390PS.INDUS     3931   LA1
# IBM390PS.MENSA     9175   ME1
# IBM390PS.NORMA     3932   LA2
/ IBM390PS.PAVO      3931   A01     MENSA
_ IBM390PS.VELA      3932   AGZ     VELA
```

Figure 5-8 IBM Z cluster list: Selecting a processor for IOCDS replace

4. The Actions on selected CPCs panel opens (see Figure 5-9). Select option `1. Work with IOCDSs`, and then press Enter.

```
+----------------- Actions on selected CPCs ------------------+
|                                                              |
|                                                              |
| Select by number or action code and press Enter.            |
|                                                              |
| 1_  1.  Work with IOCDSs . . . . . . . . . . . . . . . (s)  |
|     2.  Work with IPL attributes . . . . . . . . . . (i)    |
|     3.  Select other processor configuration . . . . (p)    |
|     4.  Work with CPC images . . . . . . . . . . . . (v)    |
|     5.  Activate HW changes remotely to a supported         |
|         processor. . . . . . . . . . . . . . . . . . (a)    |
|     6.  View current active configuration  . . . . . (c)    |
|     7.  Download active configuration  . . . . . . . (d)    |
|                                                              |
|  F1=Help     F2=Split    F3=Exit     F9=Swap    F12=Cancel  |
+--------------------------------------------------------------+
```

Figure 5-9 Actions on selected CPCs: Work with IOCDSs

5. The IOCDS List panel opens (see Figure 5-10). Select the IOCDS that you want to update for the 3931 replacement by typing a forward slash (/) next to it, and then press Enter.

```
 IOCDS List          Row 1 of 4 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select one or a group of IOCDSs, then press Enter.

                                          -----Token Match----- Write
/ IOCDS       Name     Type   Status     IOCDS/HSA IOCDS/Proc. Protect
_ A0.MENSA    IODF78   LPAR   Alternate  No        No          No
_ A1.MENSA    IODF79   LPAR   Alternate  No        No          No
_ A2.MENSA    IODF80   LPAR   POR        No        No          Yes
/ A3.MENSA    IODF77   LPAR   Alternate  No        No          No
```

Figure 5-10 IOCDS List: Selecting IOCDS for replacement

6. The Actions on selected IOCDSs panel opens (see Figure 5-11). Select option `1. Update IOCDS`, and then press Enter.

```
+---------------- Actions on selected IOCDSs -----------------+
|                                                              |
|                                                              |
| Select by number or action code and press Enter.            |
|                                                              |
| 1_  1.  Update IOCDS . . . . . . . . . . . . . (u)          |
|     2.  Switch IOCDS . . . . . . . . . . . . . (s)          |
|     3.  Enable write protection  . . . . . . (e)            |
|     4.  Disable write protection . . . . . . (w)            |
|                                                              |
|                                                              |
|                                                              |
+--------------------------------------------------------------+
```

Figure 5-11 Actions on selected IOCDSs: Update IOCDS

7. The Build IOCDSs panel opens (see Figure 5-12). Verify that all the information is correct. Complete the `Title1` field, set `Write IOCDS in preparation of upgrade` to `Yes`, and then press Enter.

```
+------------------------------- Build IOCDSs --------------------------------+
|                                                                   Row 1 of 1 |
| Command ===> ___________________________________________ Scroll ===> PAGE  |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| IODF name . . . . . . . . . : 'SYS9.IODF81'                                  |
|                                                                              |
| Title1 . IODF81_________________________________________________________     |
| Title2 : SYS9.IODF81 - 2025-04-25 09:47                                      |
|                                                                              |
|                            Write IOCDS in                                    |
| IOCDS        Switch IOCDS  preparation of upgrade                            |
| A3.MENSA     No            Yes                                               |
| ***************************** Bottom of data ****************************** |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|  F1=Help      F2=Split      F3=Exit      F5=Reset      F7=Backward           |
|  F8=Forward   F9=Swap      F12=Cancel                                        |
+-----------------------------------------------------------------------------+
```

Figure 5-12 Build IOCDSs: Verifying IODF

Tip: Specifying `Yes` in the `Write IOCDS in preparation of upgrade` field is required only when you replace or upgrade the existing hardware and want to write the IOCDS for a 9175 processor from the existing hardware. The `Yes` value enables the writing of an IOCDS that contains information that the current hardware does not recognize.

8. Because `Yes` was specified for the `Write IOCDS in preparation of upgrade` field, HCD displays a confirmation panel (see Figure 5-13). Press Enter to continue.

```
+---------- Confirm Write IOCDS in preparation of processor upgrade ----------+
|                                                                   Row 1 of 1 |
| Command ===> ___________________________________________ Scroll ===> PAGE   |
|                                                                             |
| Scroll forward to view the complete list of IOCDSs which will be written    |
| regardless of processor type in preparation of a processor upgrade. Press   |
| F3 or F12 to cancel, press ENTER to confirm the write request.              |
|                                                                             |
| The processor receiving the IOCDS(s) must be a CMOS processor.              |
|                                                                             |
| You will not be able to perform a POR using the new IOCDS until your        |
| processor has been upgraded. Do not make the new IOCDS the active one on    |
| your processor. Do not activate any I/O configuration changes in the IODF   |
| until your processor has been upgraded. Keep the old processor definition   |
| in an IODF until after the upgrade.                                         |
|                                                                             |
| IOCDS                                                                       |
| A3.MENSA                                                                    |
| ***************************** Bottom of data ****************************** |
|                                                                             |
|                                                                             |
|  F1=Help       F2=Split      F3=Exit       F7=Backward   F8=Forward        |
|  F9=Swap      F12=Cancel                                                    |
+-----------------------------------------------------------------------------+
```

Figure 5-13 Build IOCDSs: Confirm Write IOCDS

9. The Job Statement Information panel opens (see Figure 5-14). Enter the job statements as required by the installation, and press Enter. HCD submits the job to update the IOCDS.

Tip: Route the job to run on the image to which you are logged on. In that way, you know that the image can “see” the new 9175 processor to update its IOCDS.

```
+------------------------- Job Statement Information -------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the job statement information.                            |
|                                                                             |
|                                                                             |
| Job statement information                                                   |
| //WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=T                                   |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
|  F1=Help       F2=Split      F3=Exit       F5=Reset      F6=Previous        |
|  F9=Swap      F12=Cancel                                                    |
+-----------------------------------------------------------------------------+
```

Figure 5-14 Job Statement Information: Option to override job statement cards

10. Verify the job output to ensure that the IOCDS was written without error and to the correct IOCDS. You receive the following messages:

```
ICP057I IOCP JOB WIOCP    SUCCESSFUL.  LEVEL A3 IOCDS REPLACED.

Sev  Msgid    Message Text
 I CBDA674I  IOCP successfully completed for A3.MENSA.
```

11. Now, if you return to HCD option 2.11 and view the IOCDS, notice that the Systems Network Architecture (SNA) Address is still `IBM390PS.PAVO` (see Figure 5-15).

```
 Processor Cluster List                 Row 1 of 6
Command ===> _______________________________________________ Scroll ===> PAGE

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# IBM390PS.HYDRA     8562   LT2
# IBM390PS.INDUS     3931   LA1
# IBM390PS.MENSA     9175   ME1
# IBM390PS.NORMA     3932   LA2
_ IBM390PS.PAVO      3931   A01     MENSA
_ IBM390PS.VELA      3932   AGZ     VELA
```

Figure 5-15 Processor Cluster List: Selecting a processor for IOCDS verify

12. Also, when you select `IBM390PS.PAVO`, notice that IOCDS A3 (to which you wrote the upgrade IODF) has a status of `Invalid` (see Figure 5-16). This error occurs because you specified `Yes` for the `Write IOCDS in preparation for upgrade` field, and the IOCDS contains IOCP statements and code that are relevant only for a 9175 processor.

The status switches when this processor is upgraded to a 9175 processor. The 9175 IOCDS status changes to `Alternate` and the 3931 IOCDSs changes to `Invalid`.

> **Tip:** Generally, rewrite the IOCDS that is written in preparation for the upgrade at your earliest convenience. Subsequent miscellaneous equipment specifications (MESs) might cause an IOCDS that is written in preparation for an upgrade to become invalid.

```
 IOCDS List          Row 1 of 4 More:
Command ===> _______________________________________________ Scroll ===> CSR

Select one or a group of IOCDSs, then press Enter.

                                          -----Token Match----- Write
/ IOCDS       Name     Type   Status      IOCDS/HSA IOCDS/Proc. Protect
_ A0.MENSA    IODF78   LPAR   Alternate   No        No          No
_ A1.MENSA    IODF79   LPAR   Alternate   No        No          No
_ A2.MENSA    IODF80   LPAR   POR         No        No          Yes
_ A3.MENSA    IODF81   LPAR   Invalid     No        Yes         No
******************************* Bottom of data ****************************
```

Figure 5-16 IOCDS List: IOCDS verified with a status of Invalid

5.3 Creating a reset profile on the Support Element

To build and activate your reset profile by using the HMC, complete the steps in this section.

5.3.1 Background activities that occurred

The following activities are upgraded to a 9175-ME1 processor.

- A new HMC is installed with the correct driver level to support the 9175 processor, and it is connected to the customer HMC network.
- The upgraded 9175 processor (in this example MENSA) was defined to the new HMC:
  - The 3931 processor, now upgraded to a 9175 processor, underwent a POR with the Diagnostic (DEFAULT) IOCDS.
  - The 9175 processor had a new IOCP written to its IOCDS from the IODF (IODF81) by using HCD Option 2.11.
  - The 9175 processor is now ready to be customized with specific customer definitions.

5.3.2 Building the reset profile and pointing it to the required IOCDS

Now that the IOCP file is written to an IOCDS, build a reset (POR) profile to point to that IOCDS. This reset profile performs a POR for the new 9175 processor after it is upgraded and handed over from the IBM System Services Representative (IBM SSR).

To build the profile, complete the following steps:

1. Log on to the HMC workstation that is supplied with the 9175 processor with SYSPROG authority, or use a remote web browser and select the new 9175 processor.
2. Under Systems Management, click **Systems Management** to expand the list.
3. Under Systems Management, click the radio button next to the system to select it (in this example, MENSA).
4. In the Tasks window, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (see Figure 5-17 on page 102).

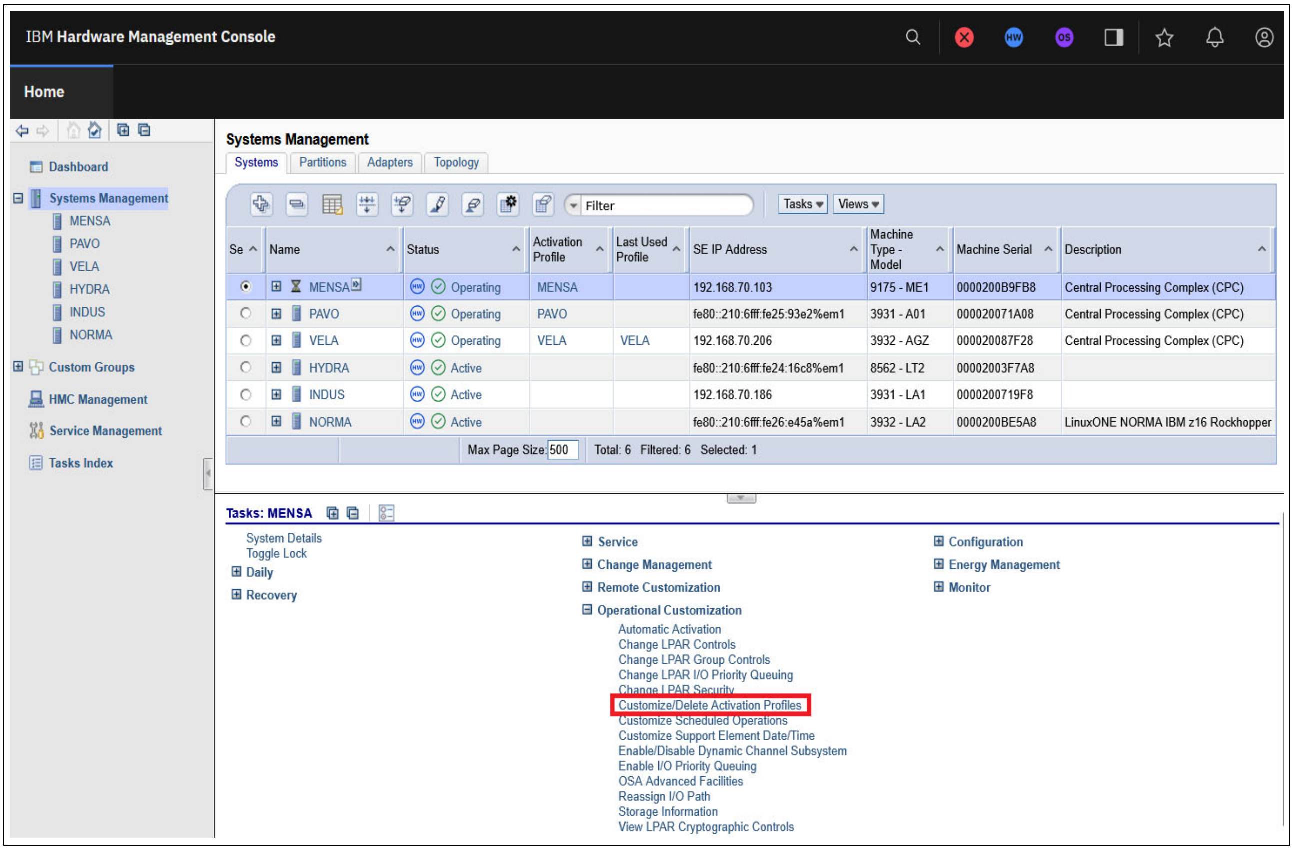

Figure 5-17 Customize/Delete Activation Profiles

5. Select the **DEFAULT** reset profile and click **Customize profile**.
6. Save this DEFAULT profile with a new profile name to be used when the POR is required (for example, TESTRESET).
7. Select the new **TESTRESET** profile and click **Customize profile**.
8. Click the IOCDS that you updated in the previous step. The ACTB0PDL message appears (see Figure 5-18).

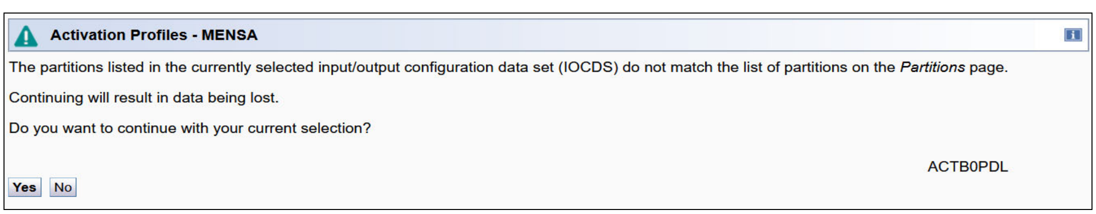

Figure 5-18 Activation Profiles: ACTB0PDL message

9. Depending on the circumstances, you can click **Yes** or **No**. You might want to review the Partition Activation List now. For this example, click **Yes**.
10. The HMC retrieves any image profiles that match the logical partition (LPAR) names that are defined in the IOCDS that was selected. You can create image profiles for those LPAR names that it cannot retrieve.

    In our example, we select **Automatically create all new images using the choices specified on this panel** and **Use the selected profile as a template when automatically creating new image profiles: DEFAULT**. Click **OK** (see Figure 5-19).

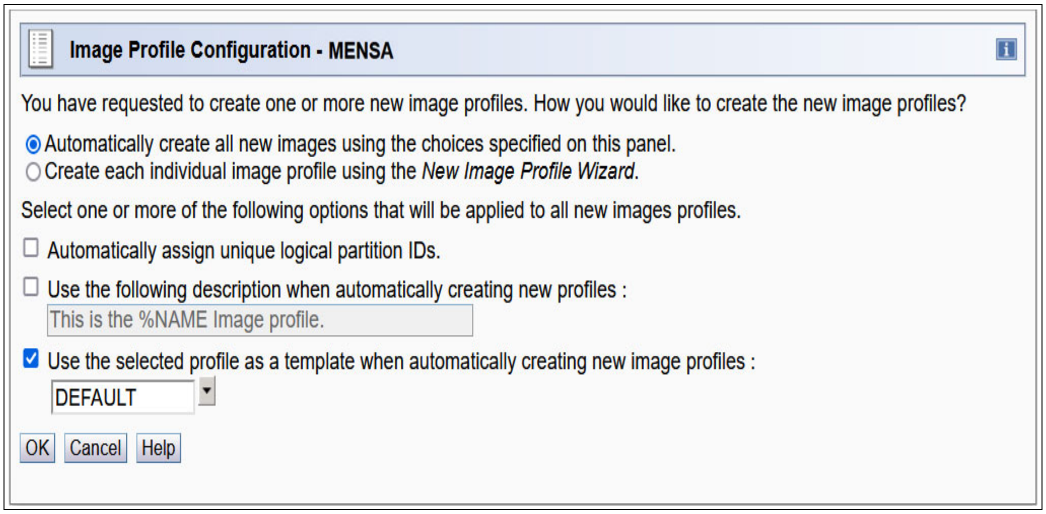

Figure 5-19 Image Profile automatic build options

11. Note the list of LPARs that were retrieved and built based on the LPARs that were defined in the selected IOCDS. Click **Save** (see Figure 5-20).

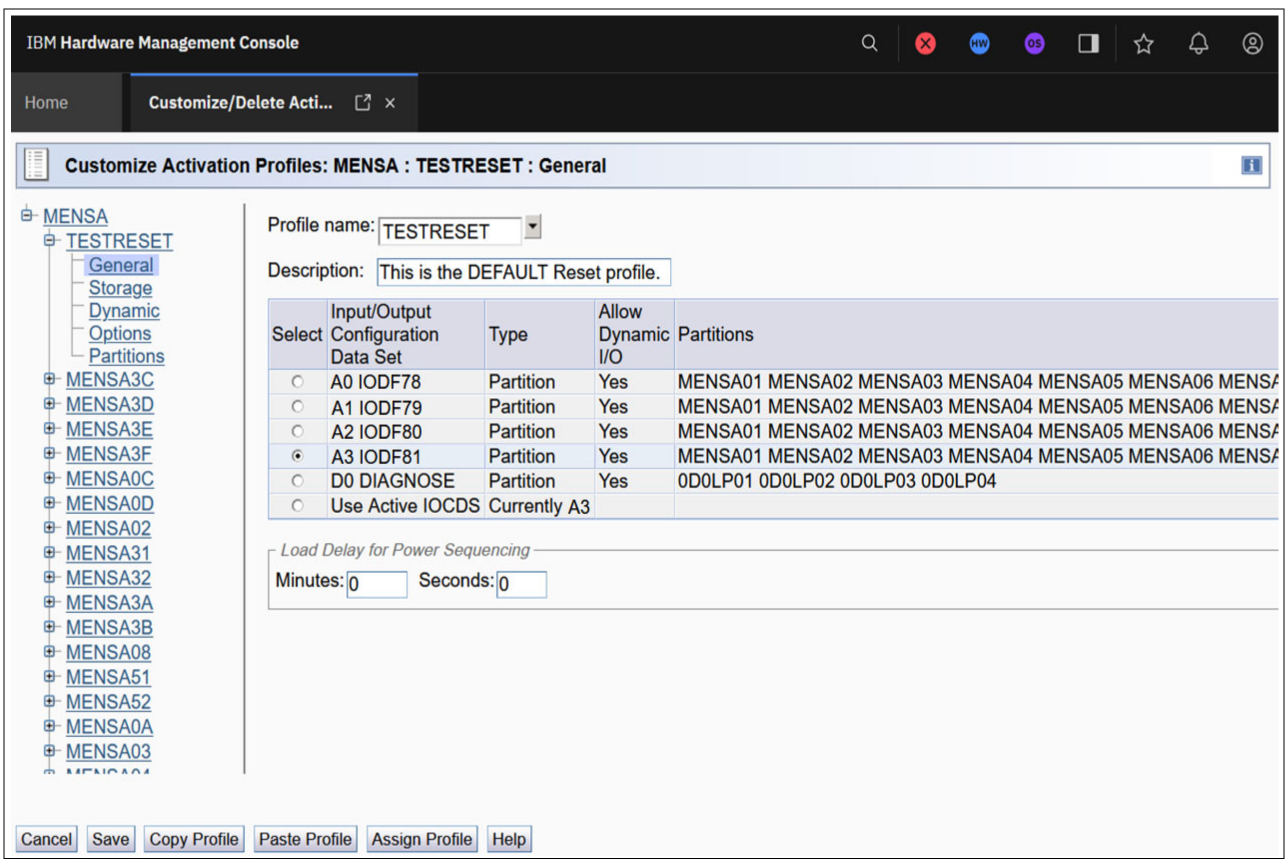

Figure 5-20 Reset and Image Profile list: Selecting IOCDS

For planning information, see Chapter 2, "Planning considerations" on page 7.

5.3.3 Setting up and verifying the reset profile

To set up and verify the reset profile, complete these steps:

1. Click **Partitions** to display the list of LPARs in the partition activation list.

   This window lists all the partitions that were retrieved by the automatic build for reset profile TESTRESET. The partition list also determines all the image profiles that would be activated if the CPC was POR.

   Here, you can tailor which image profiles are displayed and activated, and also the order of activation and the order in which they are displayed in the reset profile.

   Typing over or removing the number in the Order field determines how you want the partitions in the reset profile to behave, that is, they are removed or the order is changed.

2. After you make your determinations, click **Save** (see Figure 5-21).

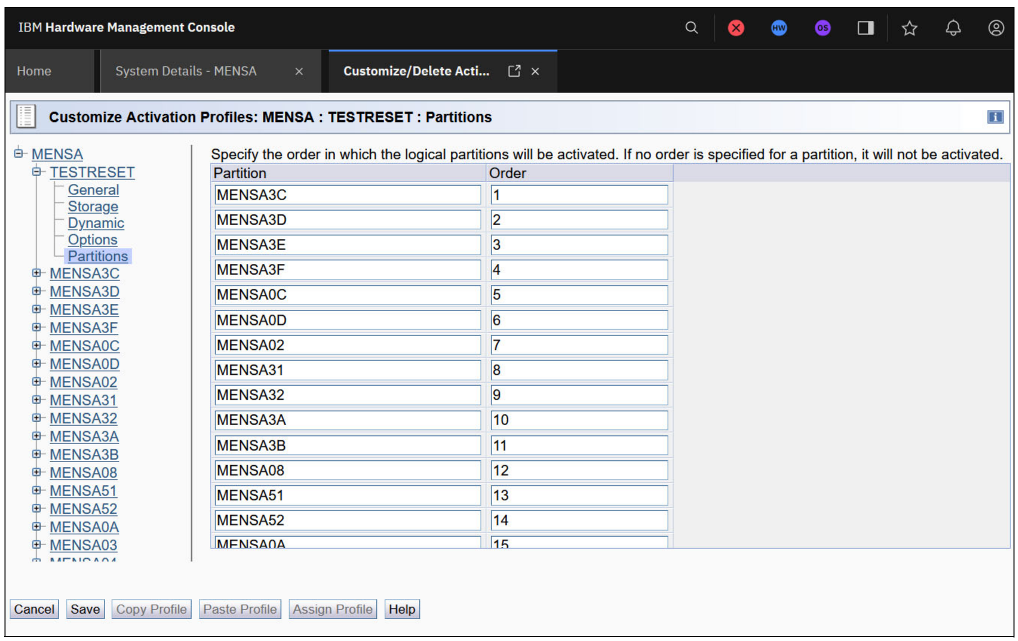

Figure 5-21 Reset and Image profile list: Updating the partition list

3. If you have any Coupling Facility (CF) partitions that are defined, HMC prompts whether you want to change the partition activation order because it is preferential but not essential that CF LPARs are activated before z/OS LPARs. Click **Yes** or **No** (see Figure 5-22).

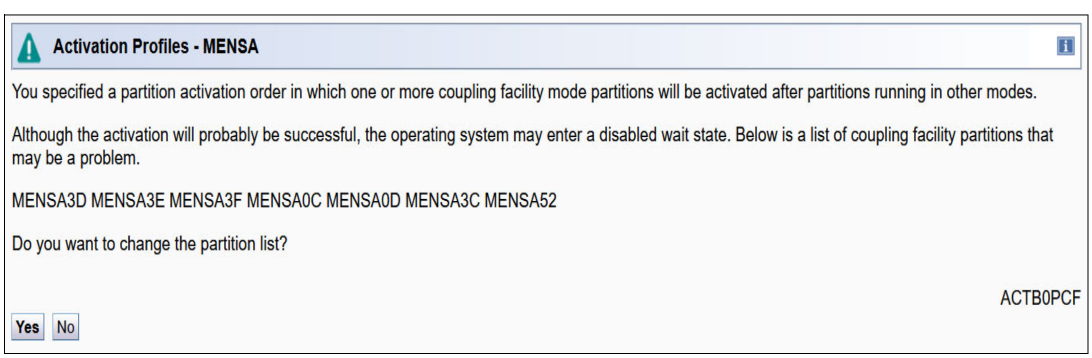

Figure 5-22 Reset and Image Profiles list: Coupling Facility LPAR verification

5.4 Creating an image profile on the Support Element

Image profiles contain all the specific parameters that relate to the partition, which include the following ones:

- General
- Processor
- Security
- Storage
- Options
- Load
- Crypto
- Time Offset

Click one of the image profiles to set up the partition parameters. In our example, we select **MENSA32**.

5.4.1 Image Profile: General page

The General page is displayed first (see Figure 5-23).

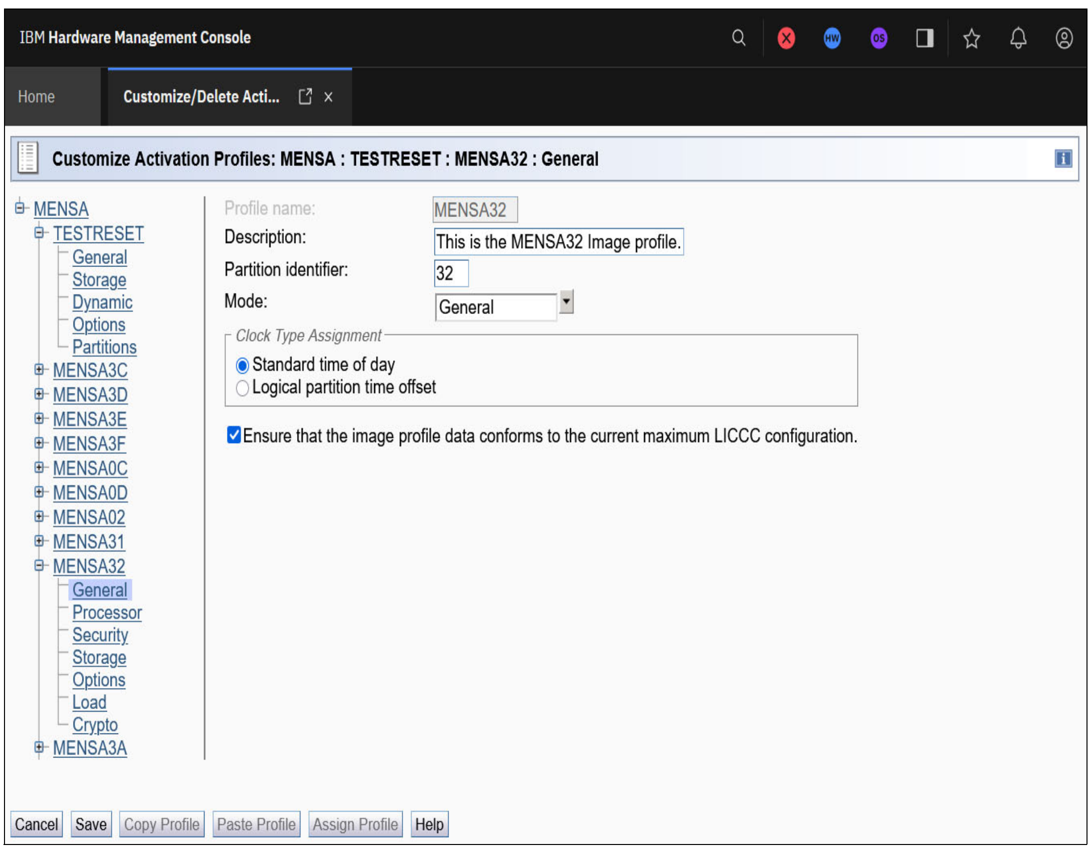

Figure 5-23 Image Profile: General

Review the following settings:

- Description
- Partition identifier
- Mode
- Clock Type Assignment:
  - Standard time of day (TOD)
  - LPAR time offset

5.4.2 Image Profile: Processor page

Click the **Processor** link to set up the partition CPU and weight information (see Figure 5-24).

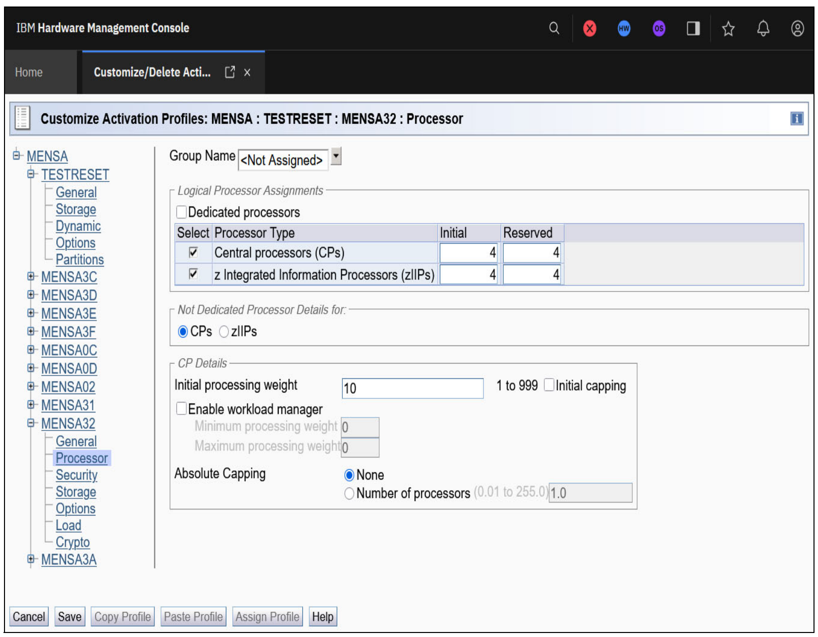

Figure 5-24 Image Profile: Processor

Review the following settings:

- Dedicated processors checkbox: Select first if you want to set dedicated central processors (CPs), IBM Z Integrated Information Processors (zIIPs), IFLs, or Internal Coupling Facilities (ICFs).
- CPs for Initial and Reserved.
- zIIPs for Initial and Reserved.
- Not dedicated Processor Details for CPs and zIIPs, IFLs, or ICFs.
- Initial processing weight.
- Initial capping.
- Enable workload manager.
- Absolute capping.

5.4.3 Image Profile: Security page

Next, click the **Security** link to set up the partition security parameters (see Figure 5-25). Review the following settings:

- Partition Security Options
- BCPii Permissions
- Counter Facility Security Options
- Sampling Facility Security Options
- CFACF Key Management Options

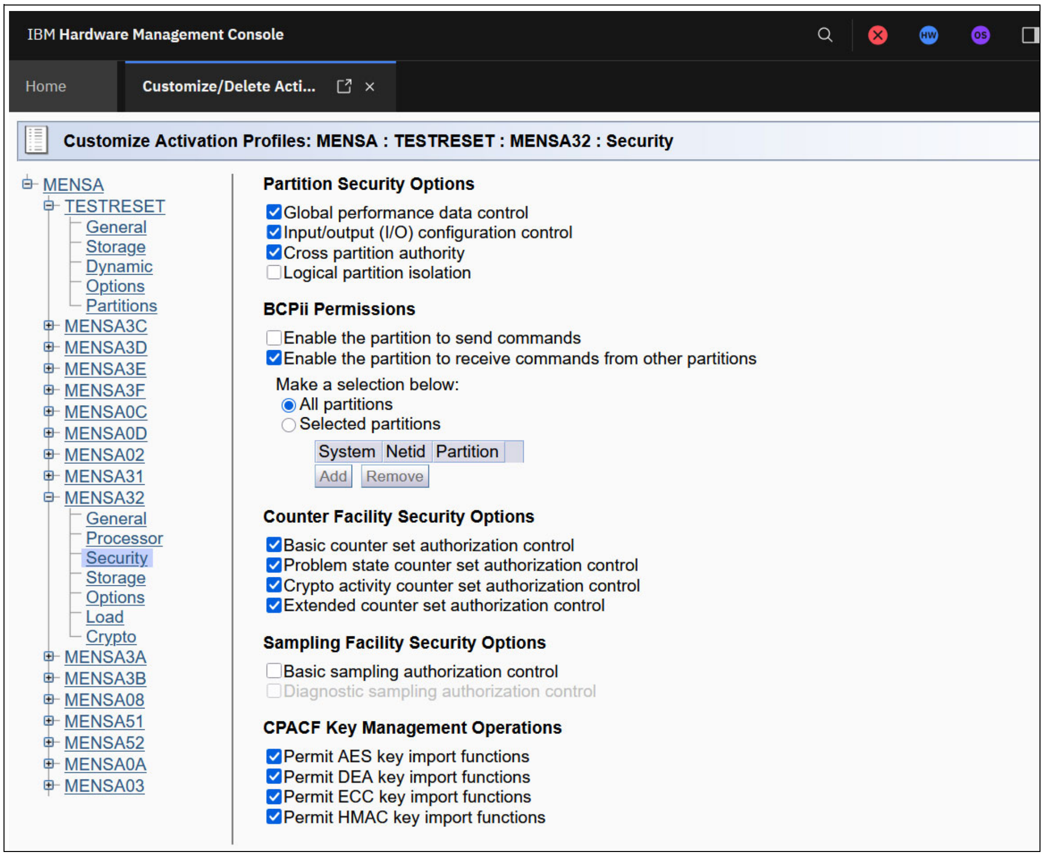

Figure 5-25 Image Profile: Security

5.4.4 Image Profile: Storage page

Click **Storage** to set up the partition Initial and Reserved storage and Virtual Flash Memory (VFM) Initial and Maximum values (see Figure 5-26). Review the following settings:

- Central Storage Amount, Initial, and Reserved
- Virtual Flash memory Initial and Maximum

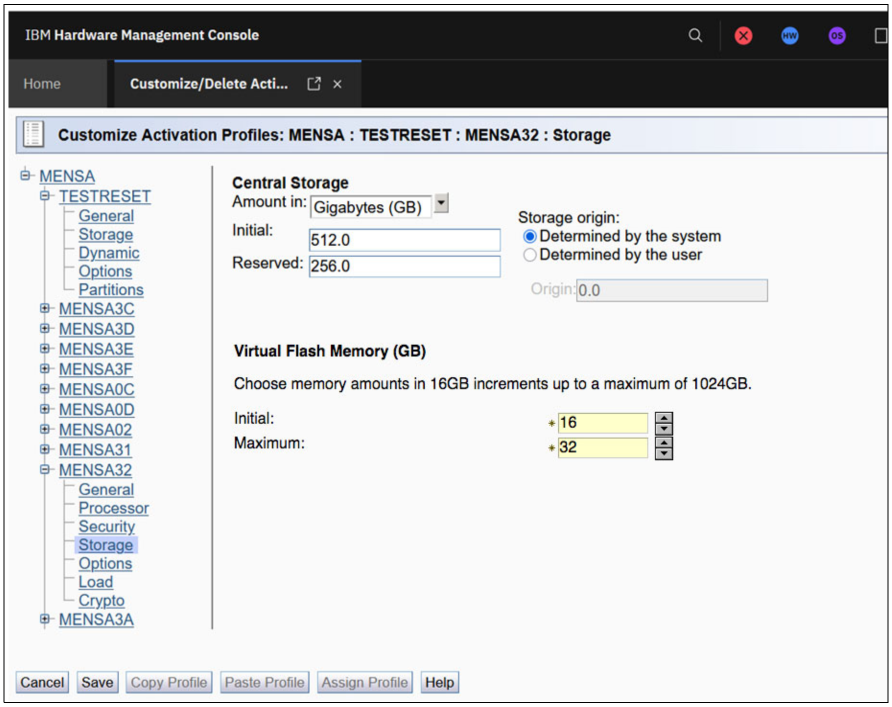

Figure 5-26 Image Profile: Storage

5.4.5 Image Profile: Options page

Click the **Options** link to set up the partitions defined capacity (see Figure 5-27).

Review the following settings:

- Minimum input/output (I/O) priority
- Maximum input/output (I/O) priority
- Defined capacity
- CP management cluster name

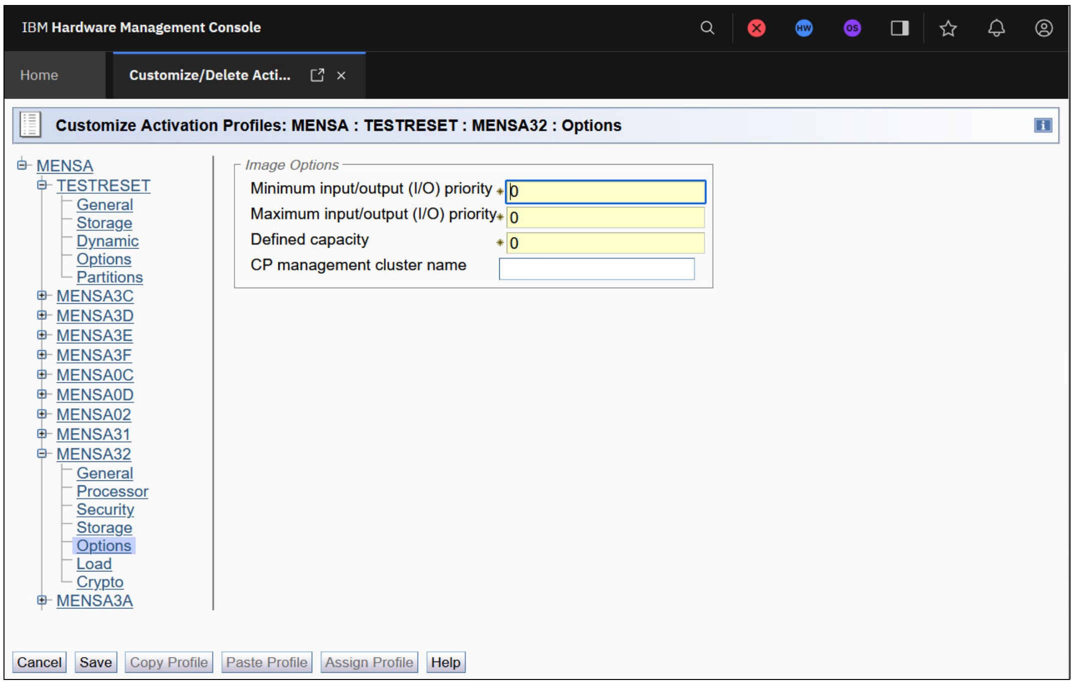

Figure 5-27 Image Profile: Options

5.4.6 Image Profile: Load page

Click the **Load** link if you want to set up any automatic load (IPL) parameters when the partition is activated by using a POR or image profile activation. If you prefer not to use IPL to load a z/OS system into a partition during a POR or image profile activation, then you can set up and activate load profiles, and then use them when they are required. This topic is mentioned in the following section (see Figure 5-28).

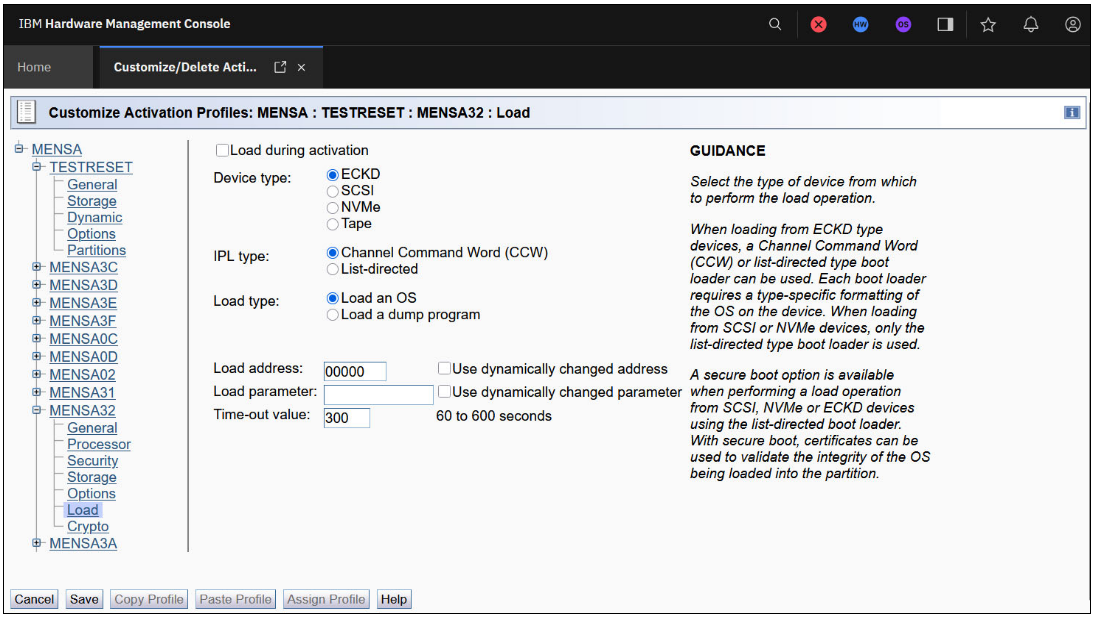

Figure 5-28 Image Profile: Load

5.4.7 Image Profile: Crypto page

Click the **Crypto** link to define the Crypto Domain Index IDs and the number of Crypto engines that are assigned to that Domain ID, and whether they are only a candidate or a candidate and online (see Figure 5-29). Review the following settings:

- Assigned Domains, which is where you first assign a Domain Index ID.
- Assigned Cryptos, which is where you assign which of and how many of the installed Crypto engines are assigned to the Domain ID and this partition.

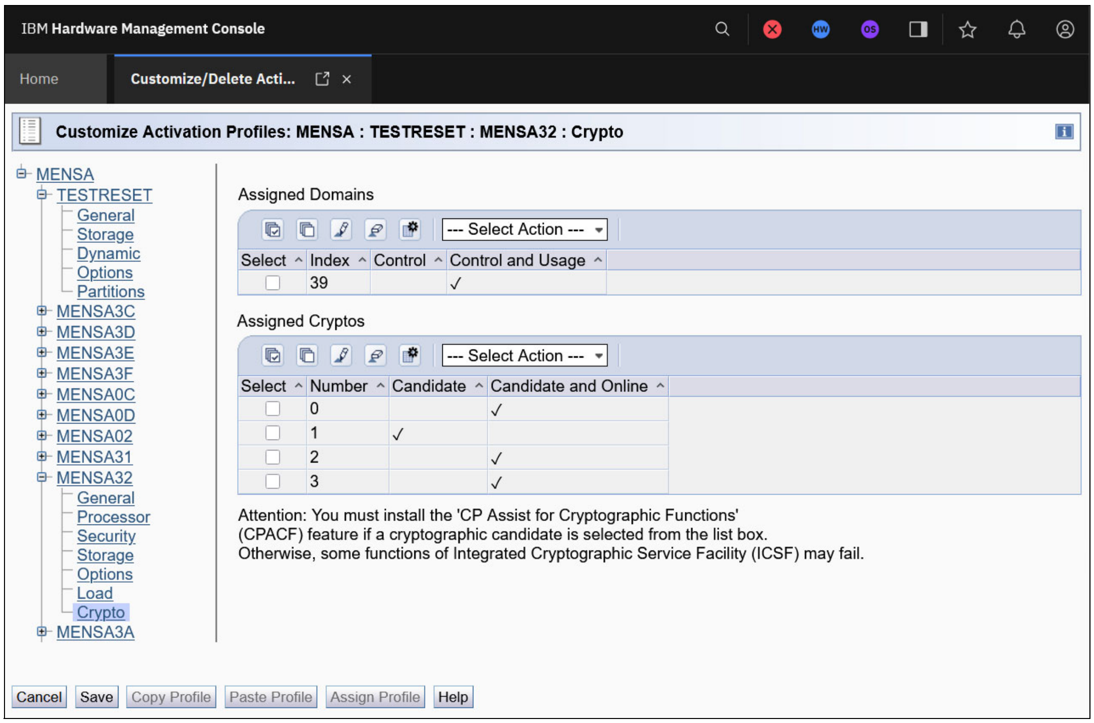

Figure 5-29 Image Profile: Crypto

5.4.8 Image Profile: Time Offset

If you selected **Logical partition time offset** in the General window, then an extra window opens in the image profile that is called Time Offset. Here, you can select the partition time offset against the CPC time as set by the Server Time Protocol (STP) (see Figure 5-30). Review the following settings:

- Offset: Days, hours, and minutes
- Decrease or Increase time value

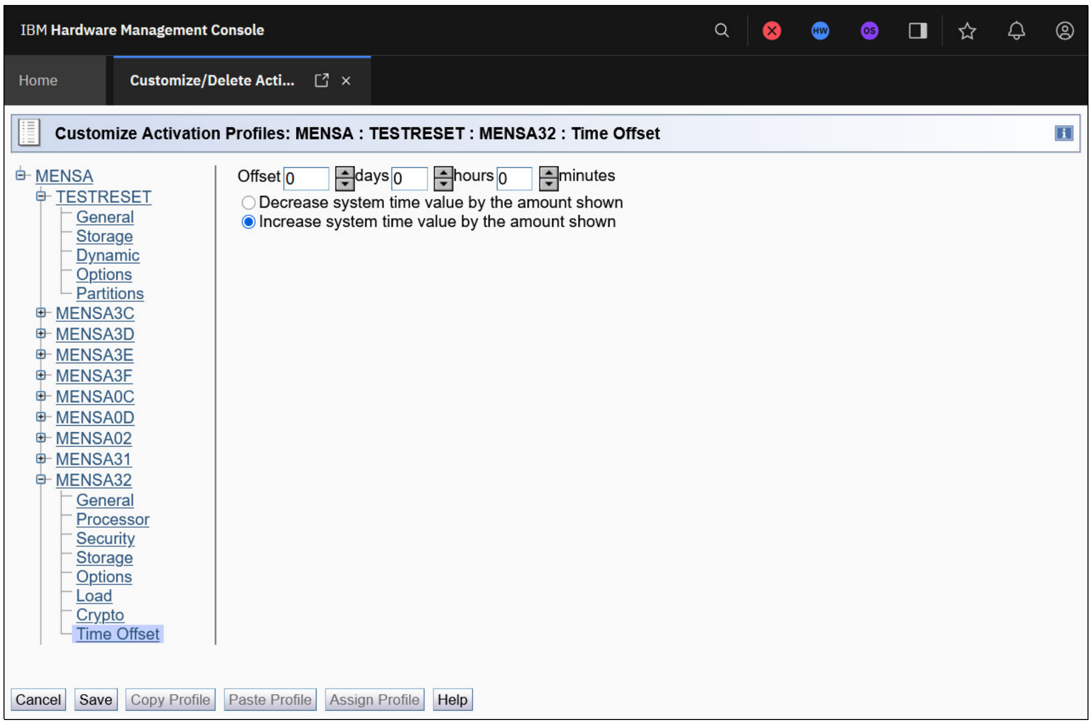

Figure 5-30 Image Profile: Time Offset

5.4.9 Image Profile: Saving

After you customize all the required image profiles for this reset profile, click **Save** to save the reset (and image) profiles for POR. If you have only a few LPARs that require activation, then perhaps it is simpler to deactivate and activate those image profiles individually. (It depends on your situation.) HMC asks for confirmation to continue to save. Click **OK** (see Figure 5-31).

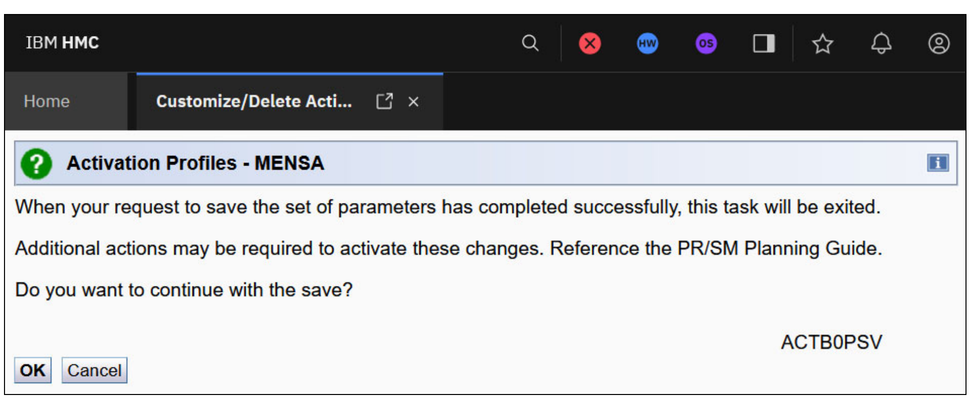

Figure 5-31 Reset and Image profile: Confirmation to save

For planning information, see Chapter 2, “Planning considerations” on page 7.

5.5 Performing a Power on Reset on the new CPC

When the 3931 processor is upgraded to a 9175 processor, your IBM SSR performs a POR with a Diagnostic IOCDS.

After this process is complete and the IBM SSR is satisfied with the status of the processor, they hand over the processor to you. Then, you can run another POR by using the reset profile that was created in 5.3.2, “Building the reset profile and pointing it to the required IOCDS” on page 101.

The 9175 processor is now ready to be activated (POR) by using the production reset profile. This process is optional but preferred depending on how many partitions that you defined on the processor.

5.5.1 Coupling Facility Links

After the POR completes with your specific customer configuration and the coupling links come online to the CF and z/OS LPARs on this CPC and any links to other CPCs, verify that they are online and established a link. One way to do this process is to display the CHPID by using Channel Problem Determination on the HMC.

To use the Channel Problem Determination process, complete the following steps:

1. Log on by using SYSPROG authority to the HMC for the new 9175 processor.
2. Click **Systems Management** to expand the list.
3. Under Systems Management, click the radio button next to the system to select it (in this example, **MENSA**).
4. In the Tasks window, click **Recovery** to expand it, and select **Single Object Operations** (see Figure 5-32).

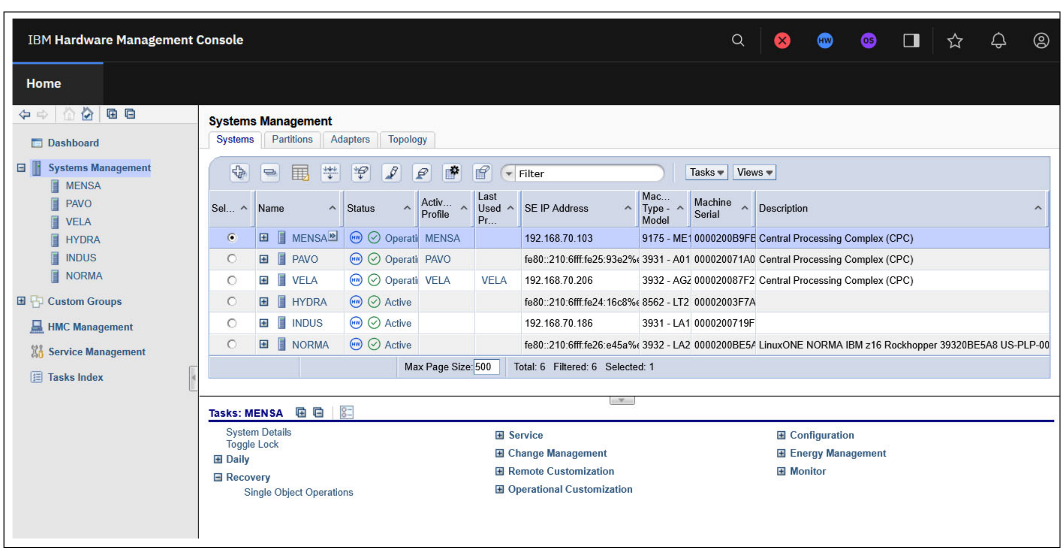

Figure 5-32 Systems Management: Main window

5. Click **OK** on the confirmation window.

6. Click **System Management** to expand the list.
7. Under Systems Management, click the CPC name to expand the options (in this example, **MENSA**).
8. Click **Partitions** to expand the list of partitions.
9. Scroll through the list of partitions until you find one of the CF partitions or z/OS partitions to which the coupling links are connected (in our example, we select **MENSA3E**).
10. Click the partition name to expand the options under the partition name.
11. Click **CHPIDs** to display the CHPID list that is specific to this LPAR (see Figure 5-33).

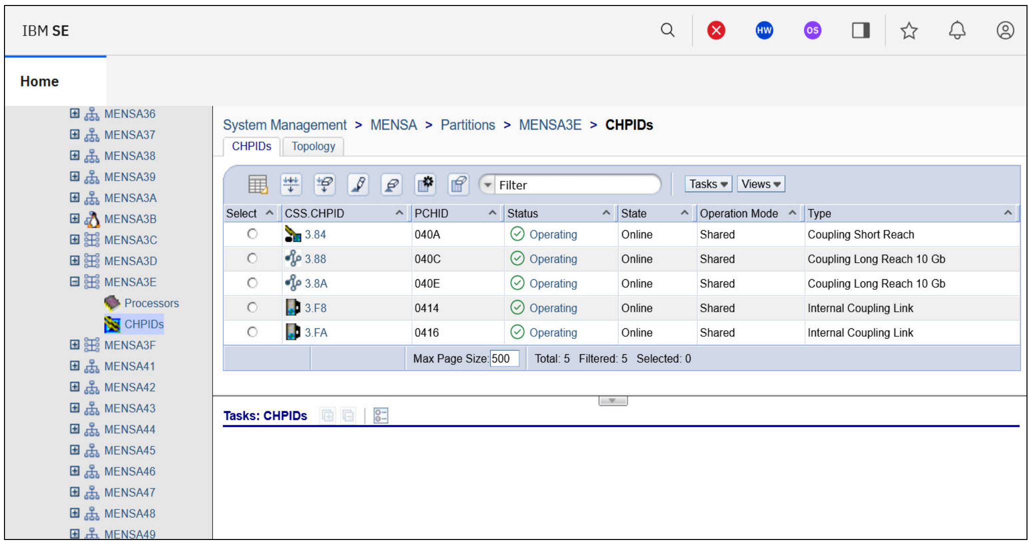

Figure 5-33 System Management: Single Object Operation

12. Select the CHPID that you want to verify. We view CHPID 3.84 (CSS=3, CHPID=84).
13. There are two ways to show the options for this CHPID: Either click the **>>** symbol next to the CHPID to expand its options, or click **CHPID Operations** to expand the options and then click **Channel Problem Determination** (see Figure 5-34).

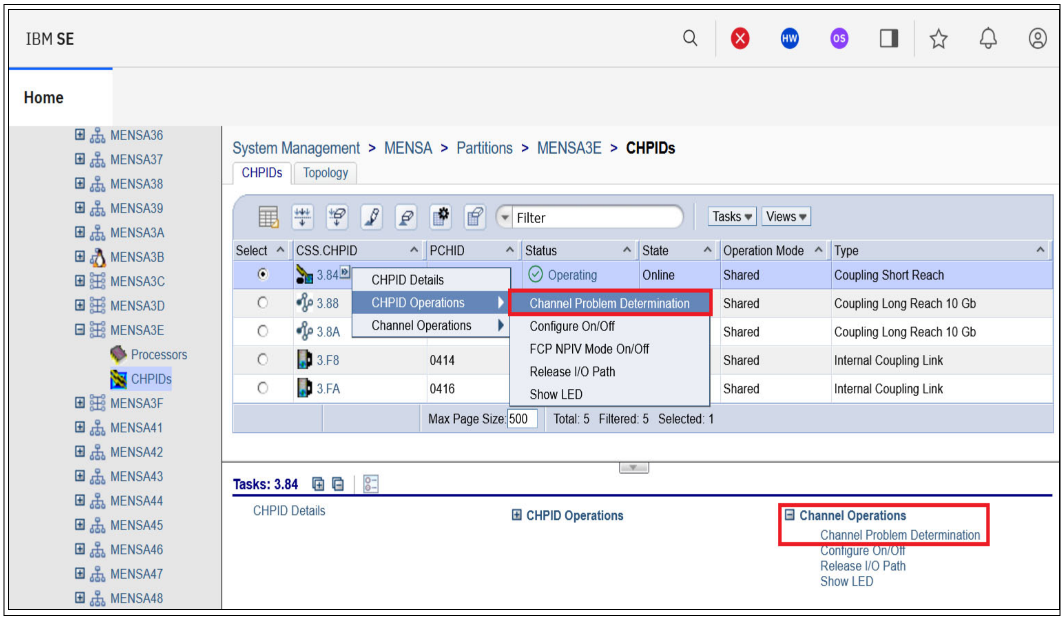

Figure 5-34 System Management: CHPID Operations

14. The HMC shows the Channel Problem Determination options. Select **Analyze channel information**, and then click **OK** (see Figure 5-35).

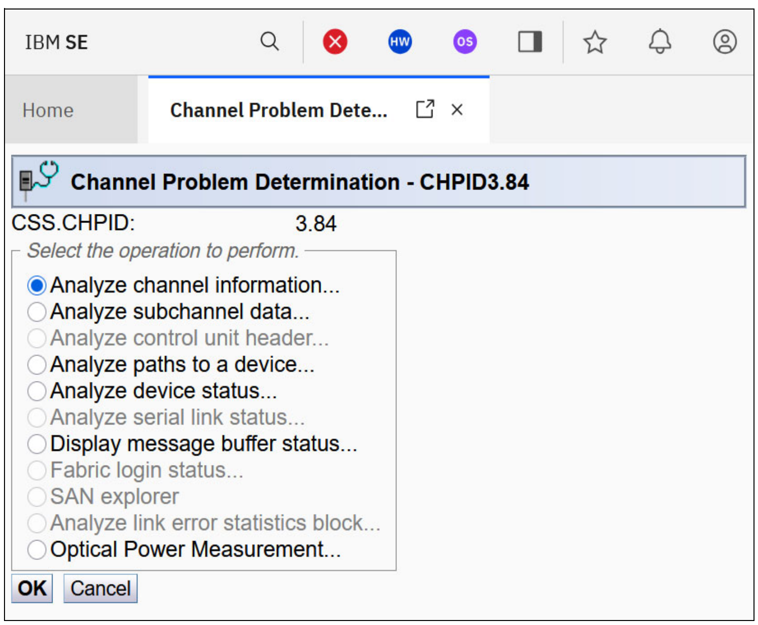

Figure 5-35 Channel Problem Determination: Analyze channel information

Note the following items (see Figure 5-36):

– State: Online
– Status: Operating
– Node type: Attached
– Node status: Valid
– Type/model: 3931-A01 (device that the CHPID is connected to)
– Seq. number: 71A08 (serial number of the device that the CHPID is connected to)
– Tag: 85 (in this case, the destination CHPID of CHPID 84)

Note the physical channel ID (PCHID) of 040A. This PCHID number is allocated by the CPC when this particular CHPID (Integrated Coupling Adapter Short Reach (ICA SR)) is defined to the HSA configuration.

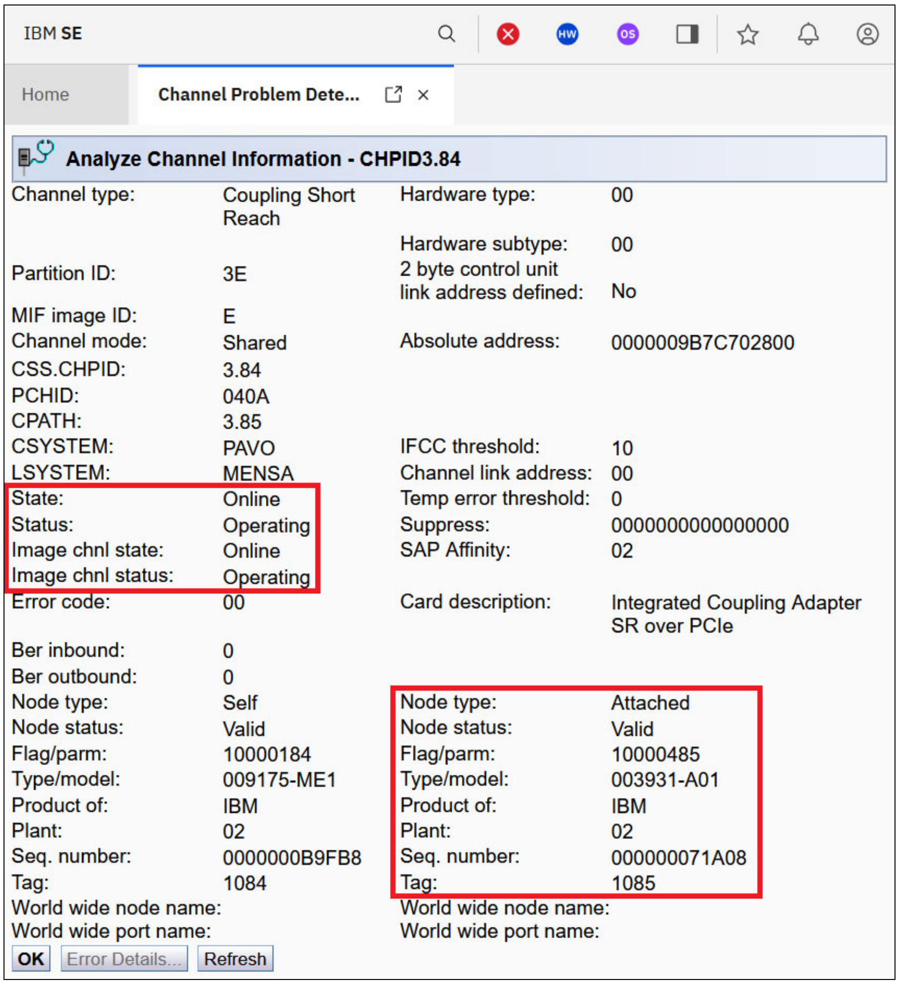

Figure 5-36 Channel Problem Determination: Display

This window verifies that the CHPID is online and operating, and also what the CHPID (cable) is connected to.

15. Continue to verify all other CF links that are defined and online.

5.5.2 Remote dynamic activation of I/O configurations for Stand Alone Coupling Facilities (SACFs), Linux on IBM Z, and z/TPF

Dynamic configuration capabilities for SACFs were added to the IBM z14 system (driver level 36) and are available on the IBM z15 and later systems.

IBM extended the Dynamic Activation of I/O configurations to Linux on IBM Z and z/TPF running on an IBM z16 CPC. This new support is applicable only when both the driving CPC and the target CPC are IBM z17 or IBM z16 with the required firmware support (Bundle S24 or higher) and when the driving system's z/OS level is 2.4 or higher with APAR OA65559.

The remote activation of dynamic changes avoids the need for disruptive hardware/firmware actions (Power-on Reset or Initial Microcode Load (IML)) to be taken to instantiate those configuration changes, reducing, or completely eliminating the client workload impact that would otherwise have resulted from taking these disruptive actions.

IBM z17 provides a supported capability to drive these hardware-only I/O configuration changes from a driving z/OS HCD instance to a remote target CPC, which is a SACF, Linux on IBM Z, and z/TPF.

Complete the following steps:

1. Provide the necessary authorization rights on your z/OS system that you use to initiate the hardware only activation. Use profiles `CBD.CPC.ACTIVATE.NetId.NAU` in class FACILTY NetId and NAU, as defined on the SE and shown in the Processor Cluster List panel:
   - READ is required for viewing and downloading the active configuration.
   - UPDATE is required for activating hardware changes only.

   > **Note:** For more information about this topic, see "Defining RACF profiles" in *z/OS HCD User's Guide*, SC34-2669.

2. Verify the RESET profile.

   Verify that the RESET profile has **Allow dynamic changes to the channel subsystem input/output definition** selected, as shown in Figure 5-37.

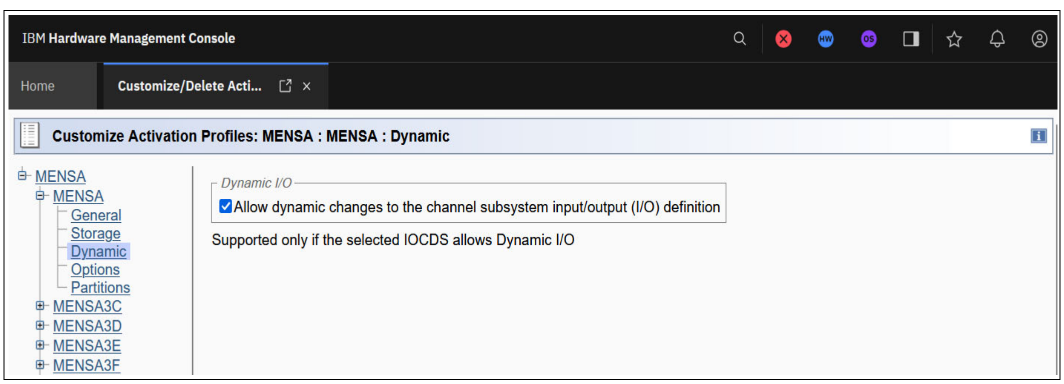

Figure 5-37 Allow dynamic changes to the channel subsystem input/output definition selected

3. After you perform the initial POR, update the reset profile by selecting **Use Active IOCDS**, which is necessary for future activations and PORs for the POR to complete by using the current IOCDS that was activated and written by using the HCD, as shown in Figure 5-38 on page 118.

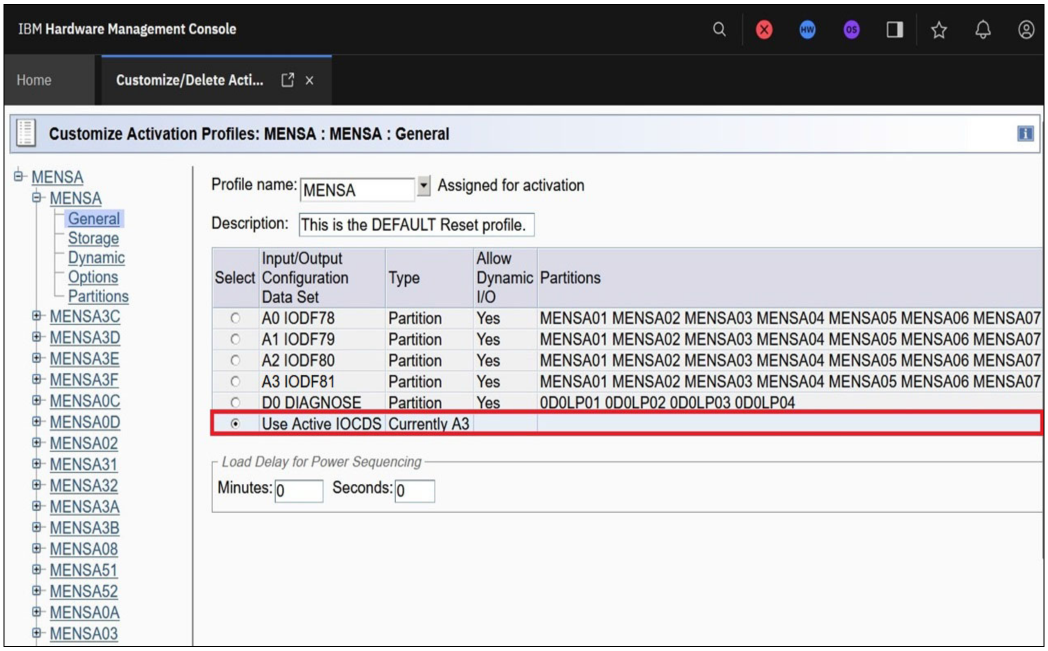

Figure 5-38 Using Active IOCDS

Note: The system is prepared for Dynamic I/O config for SACF. IBM z16 A01, IBM z16 A02, and IBM z16 AGZ or later systems support dynamic I/O config for SACF without the MCS_1 partition.

4. Make the necessary connectivity changes in your production IODF.
5. Activate the hardware configuration from an updated HCD/Hardware Configuration Manager running in z/OS LPAR on a remote CPC by selecting HCD option `2.11 Activate HW changes remotely to a supported processor (a)`, as shown in Figure 5-39.

```
Goto  Query  Help
------------- +----------------- Actions on selected CPCs ------------------+
              |                                                             |
Command ===>  |                                                             |
              | Select by number or action code and press Enter.            |
Select one or |                                                             |
              | 5_  1.  Work with IOCDSs . . . . . . . . . . . . . . (s)    |
  ----------- |     2.  Work with IPL attributes . . . . . . . . . . (i)    |
/ SNA Address |     3.  Select other processor configuration . . . . (p)    |
# IBM390PS.HY |     4.  Work with CPC images . . . . . . . . . . . . (v)    |
# IBM390PS.IN |     5.  Activate HW changes remotely to a supported         |
/ IBM390PS.ME |         processor. . . . . . . . . . . . . . . . . . (a)    |
# IBM390PS.NO |     6.  View current active configuration  . . . . . (c)    |
_ IBM390PS.PA |     7.  Download active configuration  . . . . . . . (d)    |
_ IBM390PS.VE |                                                             |
************* |                                                             | *
              +-------------------------------------------------------------+
```

Figure 5-39 Activate hardware changes only

After selecting `Activate HW changes remotely to a supported processor (a)` with no recovery required, the activation parameters are presented, as shown in Figure 5-40.

```
+------- Activate New Hardware Configuration to a supported processor -------+
|                                                                            |
|                                                                            |
| Specify or revise the values for IODF activation.                          |
|                                                                            |
| Source IODF . . . . . . : SYS9.IODF81                                      |
| Processor ID  . . . . . : MENSA                                            |
|                                                                            |
| Target IODF . . . . . . : SYS9.IODF82                                      |
| Processor ID  . . . . . . MENSA     + (press PF1 for supported processor   |
|                                       information)                         |
|                                                                            |
| Test only  . . . . . . . . . . . . . . . . . . . . . . Yes  (Yes or No)    |
| Allow hardware deletes (FORCE, FORCE=DEVICE) . . . . No   (Yes or No)      |
| Delete partition access to CHPIDs unconditionally                          |
| (FORCE=CANDIDATE)  . . . . . . . . . . . . . . . . . No   (Yes or No)      |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap       |
| F12=Cancel                                                                 |
+----------------------------------------------------------------------------+
```

Figure 5-40 Activation parameters

Note: The parameters are the same as for hardware and software (full) activation, but without a target operating system (OS) configuration and Eligible Device Table (EDT). It is a hardware-only change.

The result is presented the same way as before. After some time (the activation runs synchronously and blocks the screen), you see the results as the normal HCD Message List. It is written to the HCD message log and SYSLOG.

On the HCD panel (message list), you can see something similar to what is shown in Figure 5-41.

```
+------------------------------- Message List --------------------------------+
¦   Save  Query  Help                                                          ¦
¦  --------------------------------------------------------------------------  ¦
¦  CBDPMSGO                                                        Row 1 of 4  ¦
¦  Command ===> ___________________________________________ Scroll ===> PAGE   ¦
¦                                                                              ¦
¦  Messages are sorted by severity. Select one or more, then press Enter.     ¦
¦                                                                              ¦
¦  / Sev Msg. ID  Message Text                                                 ¦
¦  _ E   CBDA887I Following devices are to be added to processor MENSA:       ¦
¦  #              0.FF84-0.FF8F,1.FF84-1.FF8F                                  ¦
¦  _ E   CBDA889I Following control units are to be modified for processor    ¦
¦  #              MENSA: 0.FFCC,1.FFCC                                         ¦
¦  _ E   CBDD800I All change requests were successfully executed               ¦
¦  ***************************** Bottom of data ****************************** ¦
¦                                                                              ¦
¦   F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset           ¦
¦   F7=Backward   F8=Forward    F9=Swap      F10=Actions   F12=Cancel          ¦
¦  F13=Instruct   F22=Command                                                  ¦
+------------------------------------------------------------------------------+
```

Figure 5-41 Activation result

6. Similar to any other activation process, you should write a IOCDS to SACF, switch to the newly written IOCDS, and ensure that `Use Active IOCDS` is selected in the RESET profile.

7. Configure the newly added hardware to SACF and Activate New image message:

   – On the CF side, check the status by running the **`Display CHP`** command and configure channels online by running the **`Configure xx online`** command from the Operating System Messages task for the respective CF LPAR, as shown in Figure 5-42.

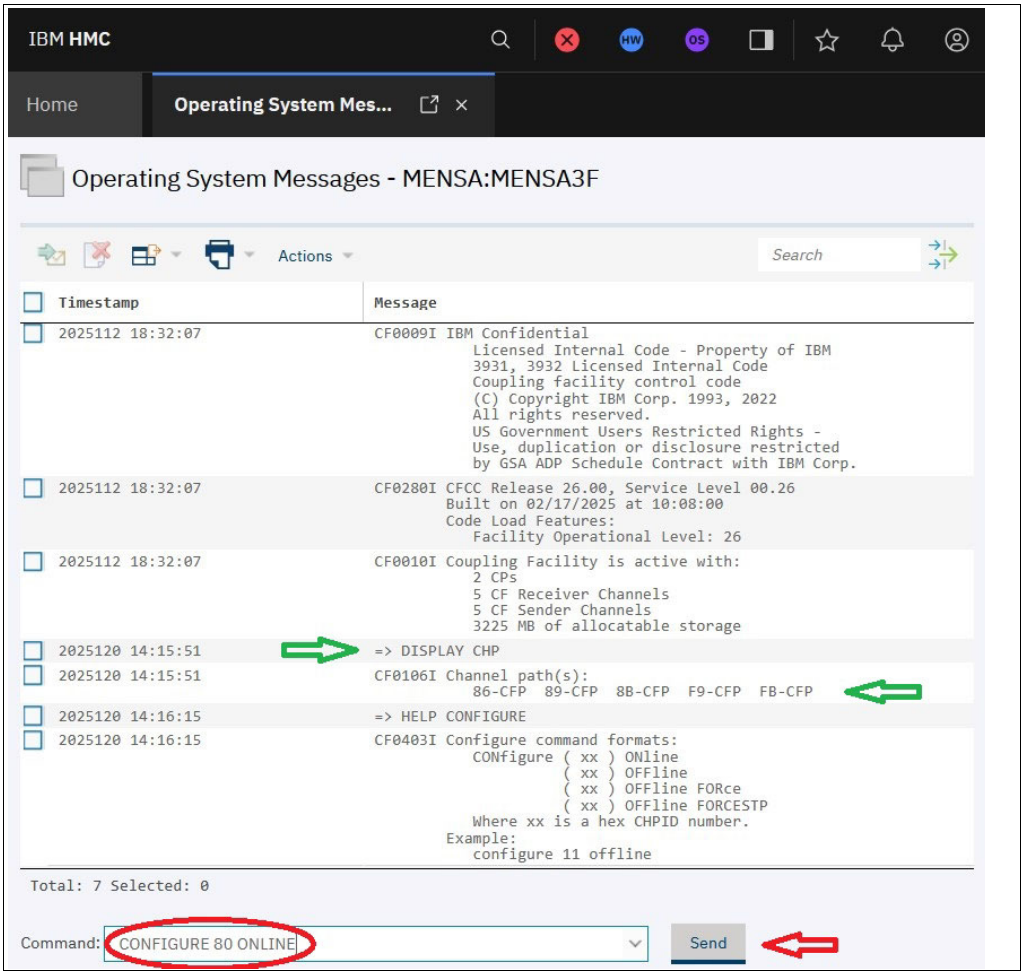

Figure 5-42 CF Operating System Message

   – On the z/OS side:

     i. Activate IODF in z/OS LPARs by using the best practice approach.

        Software changes (with VALIDATE) in all images and hardware changes in the last image per CPC. Write IOCDSs to z/OS CPCs and switch IOCDSs.

     ii. Ensure that newly added links are ONLINE to the respective z/OS LPARs.

        There are some commands that you can use to check the links, such as displaying the CHPIDs status or displaying the CF connectivity:

```
DISPLAY M=CHP
DISPLAY CF
DISPLAY CF,CFNAME=
```

        If some links are not in the expected state on the z/OS side, you can try to put them online by running the **`CF CHP(xx),online`** command.

Here are the new options on the Processor Cluster List (HCD option 2.11) for SACF activation:

- Use `View current active configuration (c)` to get information about the active configuration for the selected processor, as shown in Figure 5-43.

```
+------------------------ View Active Configuration -------------------------+
|                                                                            |
|                                                                            |
| Source IODF . . . . . . : SYS9.IODF82                                      |
| Processor ID  . . . . . : MENSA                                            |
|                                                                            |
| HSA token . . . . . . . : MENSA    25-05-05 17:58:11                        |
|                                                                            |
| Recovery required . . . : No                                               |
| Target IODF . . . . . . : SYS9.IODF82                                      |
| Processor ID  . . . . . : MENSA                                            |
|                                                                            |
|  F1=Help     F2=Split    F3=Exit     F9=Swap    F12=Cancel                 |
+----------------------------------------------------------------------------+
```

Figure 5-43 View Active Configuration

- Use Download active configuration (d) only when requested by IBM service personnel.This option shows a panel like Figure 5-44.

```
+---------------------- Download Active Configuration -----------------------+
|                                                                            |
|                                                                            |
| Selected SNA address  . : IBM390PS.MENSA                                   |
| Source IODF . . . . . . : SYS9.IODF82                                      |
|                                                                            |
| Specify the following value.                                               |
|                                                                            |
| Target IODF name  . . . . ____________________________________________  +  |
|                                                                            |
|                                                                            |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap       |
| F12=Cancel                                                                 |
+----------------------------------------------------------------------------+
```

Figure 5-44 Download Active Configuration

> **Note:** This option is available when you dynamically activate a new configuration using the steps above, but is not available when you activate a new configuration using POR.

- Use `Activate HW changes remotely to a supported processor (a)` to activate a Dynamic I/O configuration for SACF, as shown in step 5 on page 118.

5.6 Server Time Protocol configuration

Now that the CF links are verified as connected and online, you can set up the STP configuration.

The STP or Manage System Time option on the HMC under Configuration uses a GUI.

For more information about the GUI and how to set up the STP Coordinated Timing Network (CTN), see Chapter 8, “Preparing for IBM Parallel Sysplex and Server Time Protocol” on page 153.

5.7 Building and verifying load (IPL) profiles

The CPC underwent POR, the images profiles are defined and activated, the CF links are verified, and the STP and its roles are set up. Now, you can define a load (IPL) profile to use to activate (perform an IPL) an LPAR.

To build a load profile, complete the following steps:

1. Log on by using SYSPROG authority to the HMC for the new 9175 processor.
2. Under Systems Management, click **Systems Management** to expand the list.
3. Under Systems Management, select the radio button next to the system to select it (in this example, **MENSA**).
4. On the Tasks window, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (see Figure 5-45).

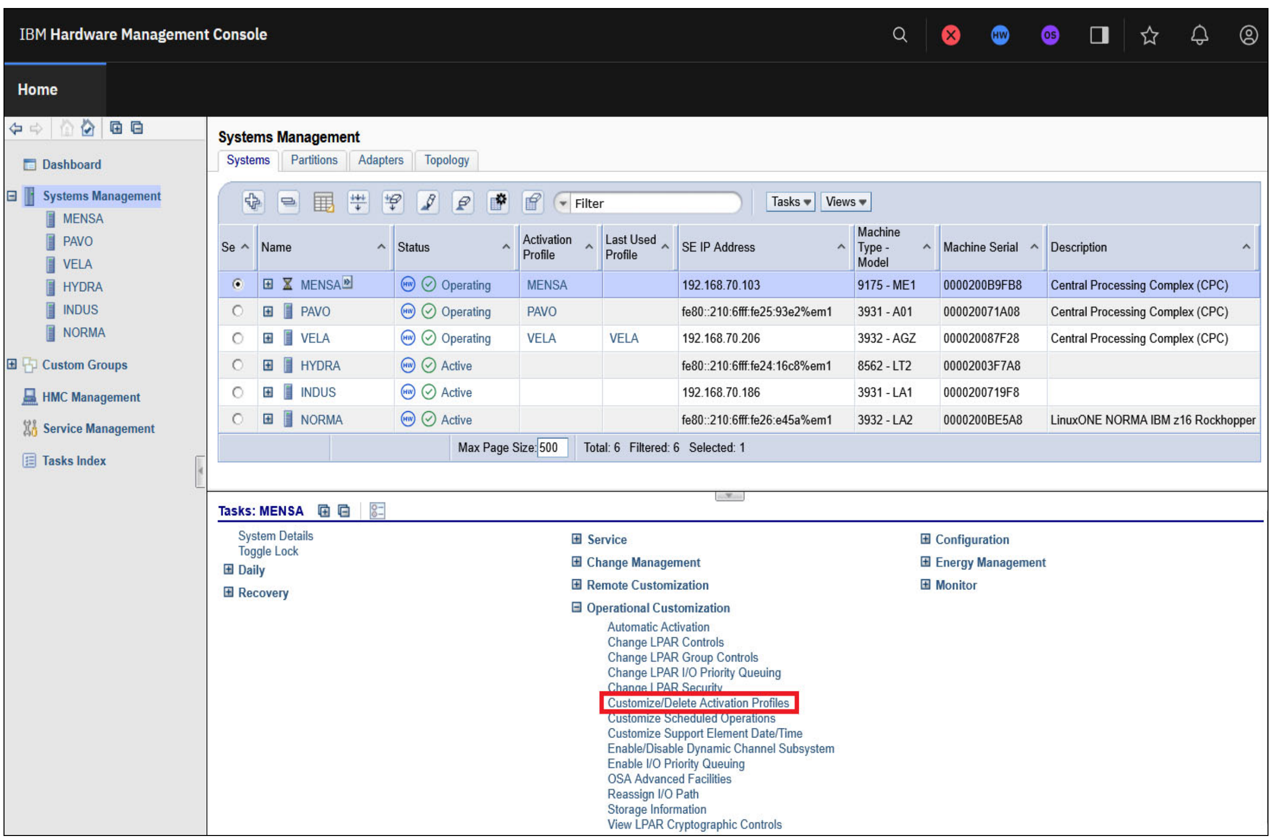

Se	Name	Status	Activation Profile	Last Used Profile	SE IP Address	Machine Type - Model	Machine Serial	Description
◉	MENSA	Operating	MENSA		192.168.70.103	9175 - ME1	0000200B9FB8	Central Processing Complex (CPC)
○	PAVO	Operating	PAVO		fe80::210:6fff:fe25:93e2%em1	3931 - A01	000020071A08	Central Processing Complex (CPC)
○	VELA	Operating	VELA	VELA	192.168.70.206	3932 - AGZ	000020087F28	Central Processing Complex (CPC)
○	HYDRA	Active			fe80::210:6fff:fe24:16c8%em1	8562 - LT2	00002003F7A8	
○	INDUS	Active			192.168.70.186	3931 - LA1	0000200719F8	
○	NORMA	Active			fe80::210:6fff:fe26:e45a%em1	3932 - LA2	0000200BE5A8	LinuxONE NORMA IBM z16 Rockhopper

Max Page Size: 500 Total: 6 Filtered: 6 Selected: 1

Figure 5-45 Systems Management: Main display

5. Select the DEFAULTLOAD load profile and click **Customize profile**. See Figure 5-46 on page 125

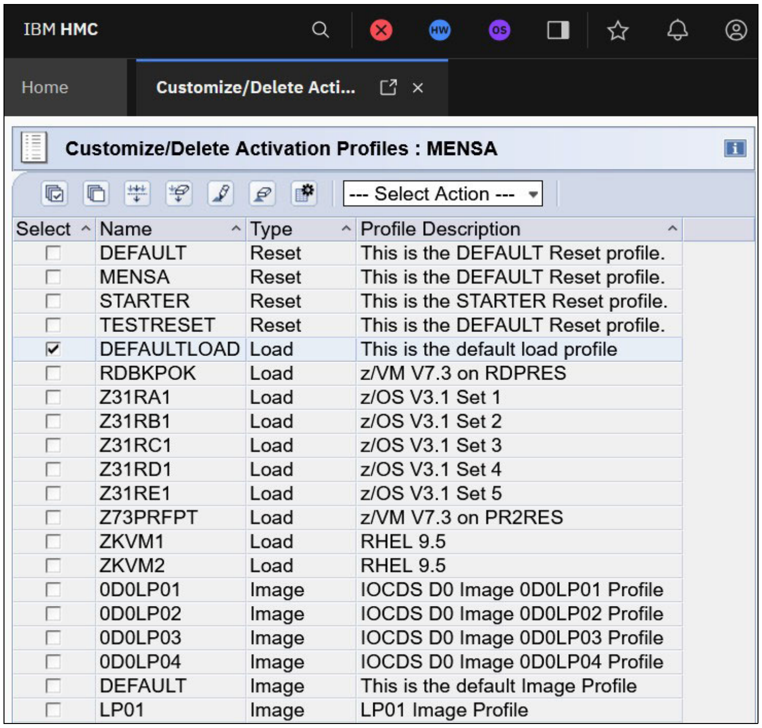

Figure 5-46 DEFAULTLOAD load profile

6. Enter the required parameters that are specific to your installation to perform an IPL:
   – Profile name: Type your preferred profile name over DEFAULTLOAD.
   – Description: Enter your preferred description.
   – Device type: ECKD.
   – IPL type: Channel Command Word (CCW).
   – Load type: Load an OS.
   – Options: Select “Clear the main memory before loading” if you clear main memory storage on the logical partition before a load.
   – Load address: The device address of the IPL volume (A075).
   – Load parameter: A04C01M1:
      • A04C: The device address of the IODF volume
      • 01: The suffix of the LOADxx member in SYS#.IPLPARM on device A04C
      • M: Automatic IPL
      • 1: SYS1.NUCLEUS

Figure 5-47 shows an example.

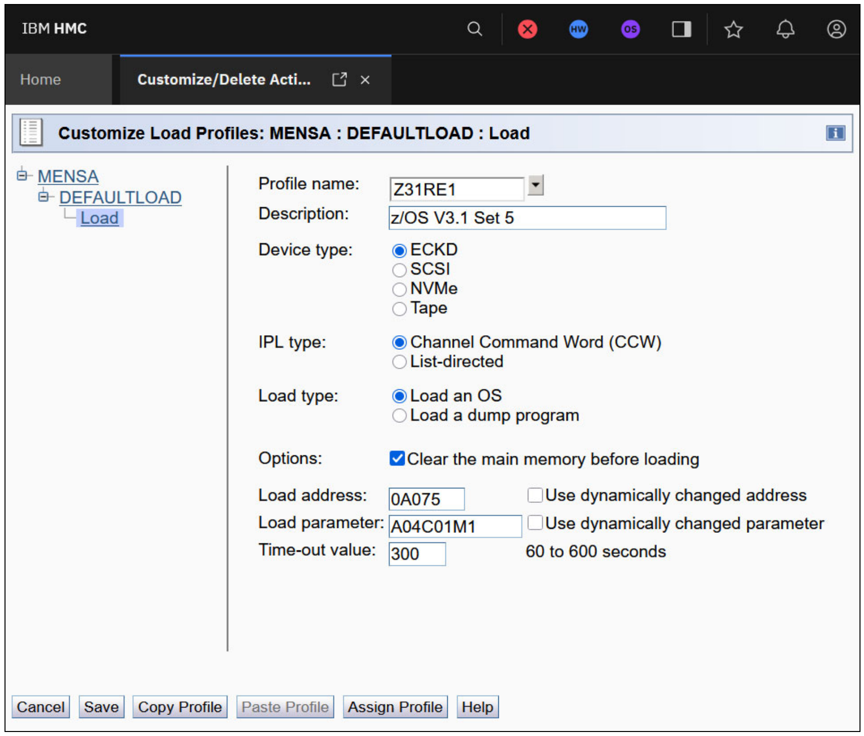

Figure 5-47 Customize Load Profiles: Load

7. Click **Save**, and then click **OK** to continue to the Save window (see Figure 5-48).

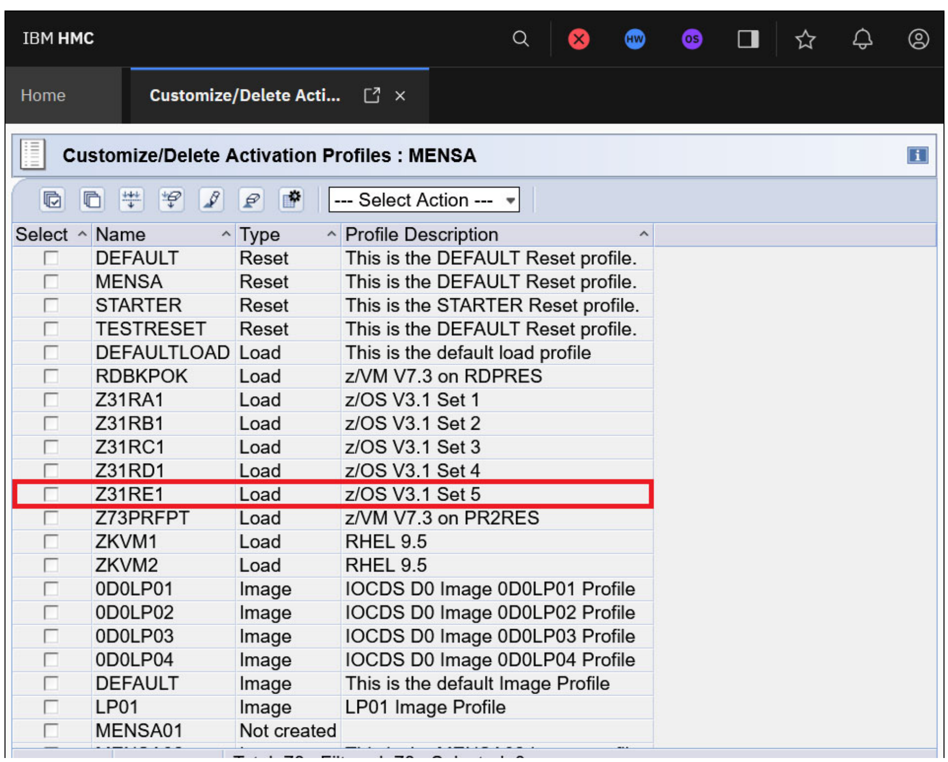

Select	Name	Type	Profile Description
	DEFAULT	Reset	This is the DEFAULT Reset profile.
	MENSA	Reset	This is the DEFAULT Reset profile.
	STARTER	Reset	This is the STARTER Reset profile.
	TESTRESET	Reset	This is the DEFAULT Reset profile.
	DEFAULTLOAD	Load	This is the default load profile
	RDBKPOK	Load	z/VM V7.3 on RDPRES
	Z31RA1	Load	z/OS V3.1 Set 1
	Z31RB1	Load	z/OS V3.1 Set 2
	Z31RC1	Load	z/OS V3.1 Set 3
	Z31RD1	Load	z/OS V3.1 Set 4
	Z31RE1	Load	z/OS V3.1 Set 5
	Z73PRFPT	Load	z/VM V7.3 on PR2RES
	ZKVM1	Load	RHEL 9.5
	ZKVM2	Load	RHEL 9.5
	0D0LP01	Image	IOCDS D0 Image 0D0LP01 Profile
	0D0LP02	Image	IOCDS D0 Image 0D0LP02 Profile
	0D0LP03	Image	IOCDS D0 Image 0D0LP03 Profile
	0D0LP04	Image	IOCDS D0 Image 0D0LP04 Profile
	DEFAULT	Image	This is the default Image Profile
	LP01	Image	LP01 Image Profile
	MENSA01	Not created	

Figure 5-48 Customize Load Profiles: New Load profile

For planning information, see Chapter 2, "Planning considerations" on page 7.

5.8 Building and verifying LOADxx members in SYS#.IPLPARM

A `LOADxx` suffix is required to perform an IPL. This data set member is stored in `SYS#.IPLPARM` on the volume that the IODF is written to. In our example, this volume is A04C (`IODFPK`). The # is the value that you use in your installation for `SYS#` data sets. The # may be any number 0 - 9, for example, `SYS0.IPLPARM`.

If you prefer to use the **`HWNAME`** keyword to point to the Processor ID, update this parameter to point to the new Processor ID (in this example, from `MENSA`). Sometimes the **`LPARNAME`** keyword is also used in the `LOADxx` members, and it might need to be reviewed or updated, such as `MENSA32`.

> **Note:** If you are going to share a `LOADxx` member with many partitions, then the **`HWNAME`** and **`LPARNAME`** keywords are required.

To build and verify `LOADxx` members in `SYS#.IPLPARM`, complete the following steps:

1. Log on by using Time Sharing Option (TSO) to a system that has access to the `SYS#.IPLPARM` data set that is on the IODF volume that you use to perform the IPL.
2. Edit data set `SYS#.IPLPARM` and edit member `LOADxx`. Figure 5-49 shows the settings that are used in our example.

```
File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
-------------------------------------------------------------------------------
EDIT       SYS0.IPLPARM(LOAD01) - 01.99                   Columns 00001 00072
Command ===>                                                 Scroll ===> CSR
000015 *-----------DEFINITION FOR SC75-------------*
000016 HWNAME   VELA
000017 LPARNAME VELA23
000018 SYSPLEX  PLEX75   Y
000019 IODF     ** SYS9     ITSO     01 Y
000020 SYSCAT   BH5CAT123CMCAT.BH5CAT
000021 NUCLST   4A
000022 PARMLIB  SYS1.PARMLIB
000023 PARMLIB  SYS1.IBM.PARMLIB
000024 PROCVIEW CORE,CPU_OK
000025 *-----------
000026 HWNAME   MENSA
000027 LPARNAME MENSA32
000028 SYSPLEX  PLEX75   Y
000029 IODF     ** SYS9     ITSO     01 Y
000030 SYSCAT   BH5CAT123CMCAT.BH5CAT
000031 NUCLST   4A
000032 PARMLIB  SYS1.PARMLIB
000033 PARMLIB  SYS1.IBM.PARMLIB
000034 PROCVIEW CORE,CPU_OK
```

Figure 5-49 z/OS: SYS#.IPLPARM: LOADxx member

- The **HWNAME** keyword is set to MENSA.
- The **LPARNAME** keyword is set to MENSA32.
- The **IODF** keyword is set to ** (where ** directs the IPL to look at what IODF / IOCP underwent a POR into the CPCs HSA, and then look on the IODF volume for that corresponding IODF). A specific IODF suffix number can be defined in the LOADxx member if you must override the HSA match.
- The **IODF** keyword points to the High-Level Qualifier of the IODF data set (SYS9) and the operating system configuration (OSCONFIG) that this system uses (ITSO). The OSCONFIG is the other part of an IODF that the IPL uses to determine which devices it can access, along with Nucleus Initialization Program Consoles and Esoterics.
- The remaining parameters are used for z/OS and not for the IODF.

5.9 Communicating information about the new CPC

Now that you made a new CPC with a new name in the configuration, you might want to communicate the new configuration specifics to the operations and support community in your organization.

6

Configuring network features

This chapter explains how to configure the network features of your IBM z17 ME1 system to connect to your network.

Naming: This publication covers the IBM z17 ME1. Throughout this chapter, we might refer to these machines as IBM z17.

This chapter includes the following topics:

- Configuring OSA for OSA-Express and for Network Express Adapter
- Configuring Shared Memory Communications (SMC-D and SMC-R)

Note: The CHPID type OSE is not supported for IBM z17, neither on OSA-Express nor on Network Express.

The former RoCE cards of previous machines IBM z15 (FC0411) and IBM z16 (FC0412, 0430, 0432, 0440, 0441, 0443, 0450, 0452) are not supported on IBM z17.

The only adapter type providing RoCE functionality for IBM z17 is the Network Express Adapter with the NETH PCIe function described in Chapter 15, “Adding Peripheral Component Interconnect Express devices” on page 331.

6.1 Configuring OSA for OSA-Express and Network Express Adapter

OSA communications can be defined either on an OSA-Express7S (see “OSA-Express” on page 34 for a complete list of supported OSA-Express features) with channel type OSD, or on a Network Express Adapter with channel type OSH. Besides from this, the I/O definition of OSA is the same for both types of adapters.

To configure OSA, the CHPID, Control Unit and devices need to be defined to the I/O configuration using HCD. This is explained in detail in Chapter 13, “Adding network devices” on page 265.

Customization of OSA for either OSD or OSH can then be performed using OSA Advanced Facilities on the HMC.

OSA Advanced Facilities is a tool that is integrated into the Hardware Management Console (HMC). To start OSA Advanced Facilities, log in to the HMC with the proper authority, and select the central processor complex (CPC) that requires OSA customization. Then, select **Operational Customization** → **OSA Advanced Facilities** (Figure 6-1).

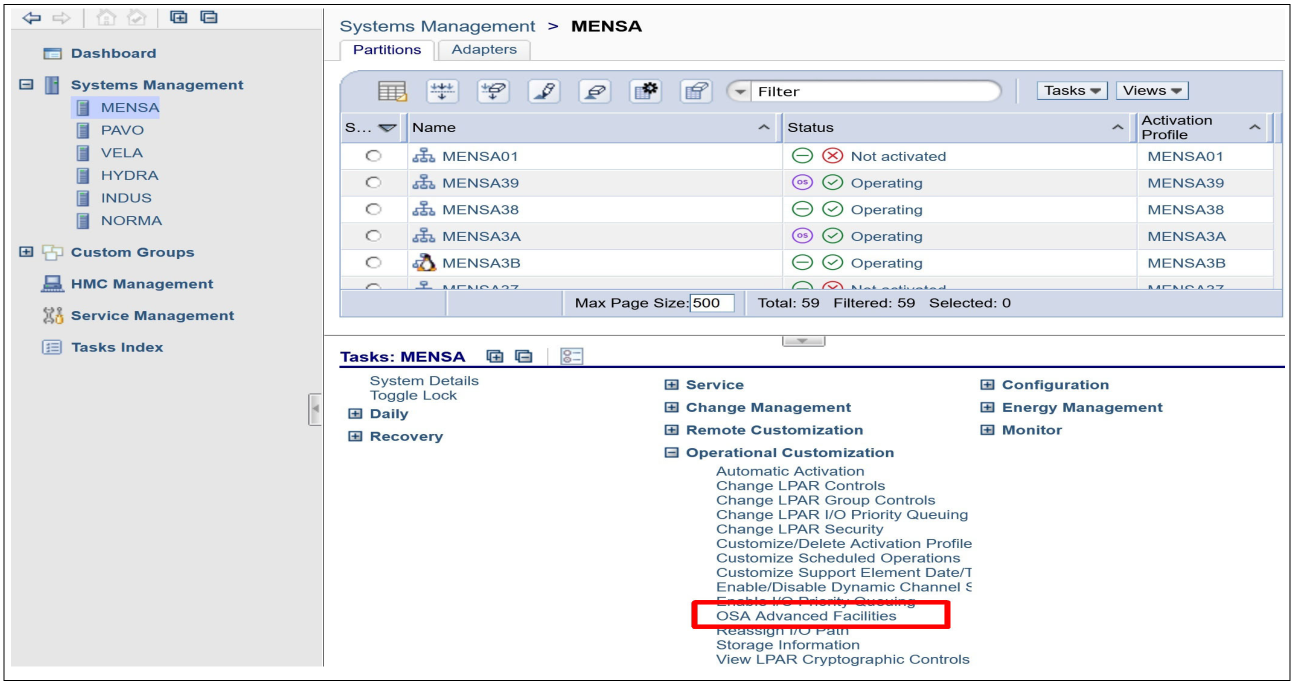

Figure 6-1 Selecting OSA Advanced Facilities from CPC

The OSA Advanced Facilities window that is shown in Figure 6-2 opens.

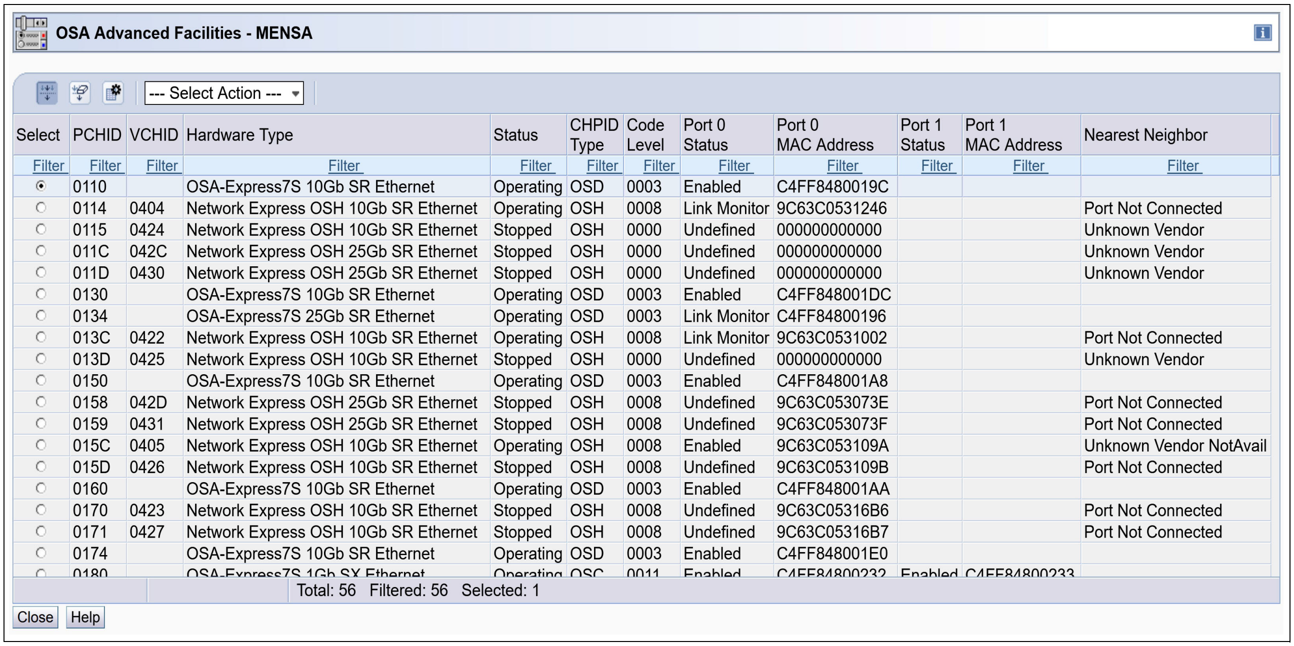

Select	PCHID	VCHID	Hardware Type	Status	CHPID Type	Code Level	Port 0 Status	Port 0 MAC Address	Port 1 Status	Port 1 MAC Address	Nearest Neighbor
Filter	Filter	Filter	Filter	Filter	Filter	Filter	Filter	Filter	Filter	Filter	Filter
◉	0110		OSA-Express7S 10Gb SR Ethernet	Operating	OSD	0003	Enabled	C4FF8480019C			
○	0114	0404	Network Express OSH 10Gb SR Ethernet	Operating	OSH	0008	Link Monitor	9C63C0531246			Port Not Connected
○	0115	0424	Network Express OSH 10Gb SR Ethernet	Stopped	OSH	0000	Undefined	000000000000			Unknown Vendor
○	011C	042C	Network Express OSH 25Gb SR Ethernet	Stopped	OSH	0000	Undefined	000000000000			Unknown Vendor
○	011D	0430	Network Express OSH 25Gb SR Ethernet	Stopped	OSH	0000	Undefined	000000000000			Unknown Vendor
○	0130		OSA-Express7S 10Gb SR Ethernet	Operating	OSD	0003	Enabled	C4FF848001DC			
○	0134		OSA-Express7S 25Gb SR Ethernet	Operating	OSD	0003	Link Monitor	C4FF84800196			
○	013C	0422	Network Express OSH 10Gb SR Ethernet	Operating	OSH	0008	Link Monitor	9C63C0531002			Port Not Connected
○	013D	0425	Network Express OSH 10Gb SR Ethernet	Stopped	OSH	0000	Undefined	000000000000			Unknown Vendor
○	0150		OSA-Express7S 10Gb SR Ethernet	Operating	OSD	0003	Enabled	C4FF848001A8			
○	0158	042D	Network Express OSH 25Gb SR Ethernet	Stopped	OSH	0008	Undefined	9C63C053073E			Port Not Connected
○	0159	0431	Network Express OSH 25Gb SR Ethernet	Stopped	OSH	0008	Undefined	9C63C053073F			Port Not Connected
○	015C	0405	Network Express OSH 10Gb SR Ethernet	Operating	OSH	0008	Enabled	9C63C053109A			Unknown Vendor NotAvail
○	015D	0426	Network Express OSH 10Gb SR Ethernet	Stopped	OSH	0008	Undefined	9C63C053109B			Port Not Connected
○	0160		OSA-Express7S 10Gb SR Ethernet	Operating	OSD	0003	Enabled	C4FF848001AA			
○	0170	0423	Network Express OSH 10Gb SR Ethernet	Stopped	OSH	0008	Undefined	9C63C05316B6			Port Not Connected
○	0171	0427	Network Express OSH 10Gb SR Ethernet	Stopped	OSH	0008	Undefined	9C63C05316B7			Port Not Connected
○	0174		OSA-Express7S 10Gb SR Ethernet	Operating	OSD	0003	Enabled	C4FF848001E0			
○	0180		OSA-Express7S 1Gb SX Ethernet	Operating	OSC	0011	Enabled	C4FF84800232	Enabled	C4FF84800233	

Figure 6-2 OSA Advanced Facilities initial window

This chapter only handles OSH and OSD channels. In the OSA Advanced Facilities initial window (Figure 6-2) the channel types OSC and NETH are also listed. The OSC type channels are covered in Chapter 7, "Defining console communication" on page 137, and the NETH function is covered in Chapter 15, "Adding Peripheral Component Interconnect Express devices" on page 331.

6.1.1 Open Card Specific Advanced Facilities

From the initial window you can select the OSA-Express (OSD) or Network Express Adapter (OSH) PCHID you want to customize and then select the action "Card Specific Advanced Facilities" from the pull down menu at the top of the window like shown in Figure 6-3.

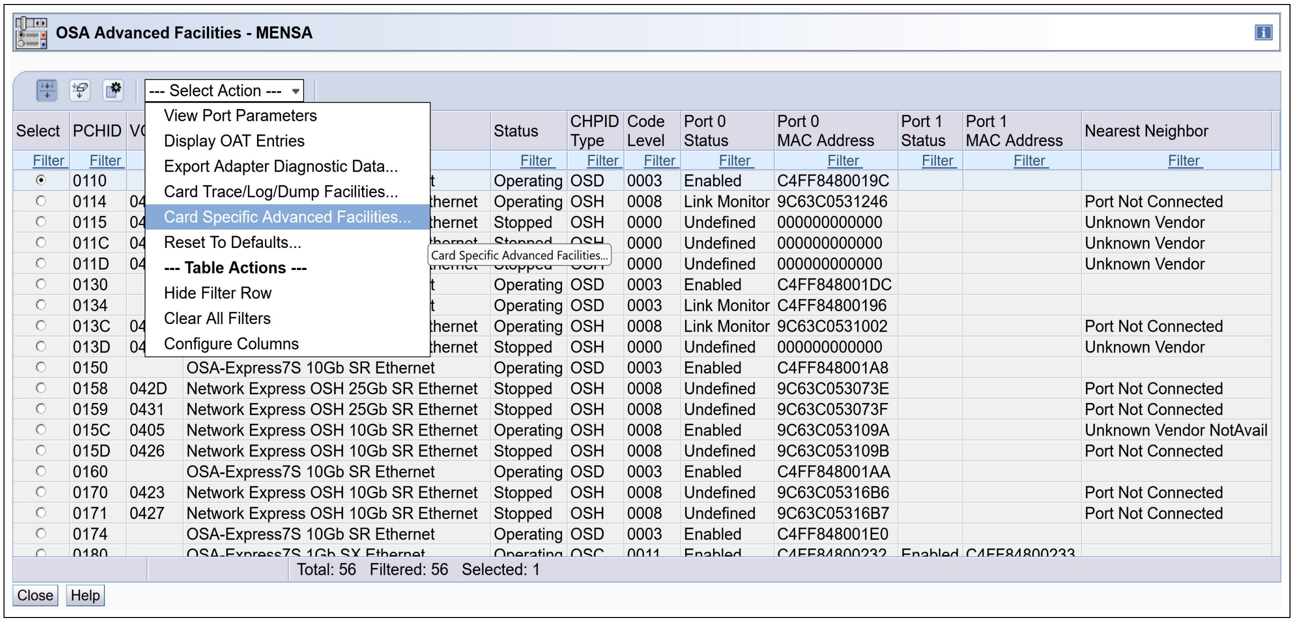

Figure 6-3 Select Card Specific Advanced Facilities

Depending of the selected channel type, a selection window will come up that allows you to choose out of some or all of the following configuration options:

- Display or alter MAC address (OSD and OSH)
- Enable or disable ports (OSD and OSH)
- Set card mode (only OSD; LAN port type 1000Base-T Ethernet)

6.1.2 Display or alter MAC address

The configuration option "Display or alter MAC address" is available for OSD and OSH channels.

After selecting this option and pressing the **OK** button, the Display or alter MAC address window opens (Figure 6-4). Set the MAC address that you want and then click **OK**.

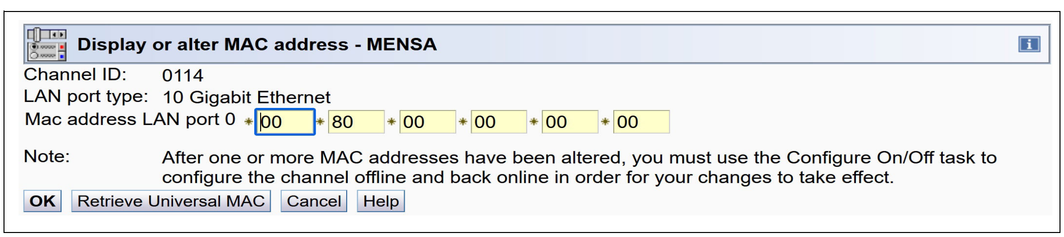

Figure 6-4 OSA/SF on the HMC: Display or alter MAC address value

For OSA-Express Adapters with two ports, the MAC addresses of both ports will be shown in the window, and you will be able to change MAC addresses of both ports at the same time.

To activate the modification, you have to configure the CHPID OFFLINE and ONLINE from each LPAR where the CHPID is defined.

6.1.3 Enable or disable ports

The configuration option “Enable or disable ports” is available for OSD and OSH channels.

After selecting this option and pressing the **OK** button, the panel ”Enable or disable ports” opens (Figure 6-5 on page 133).

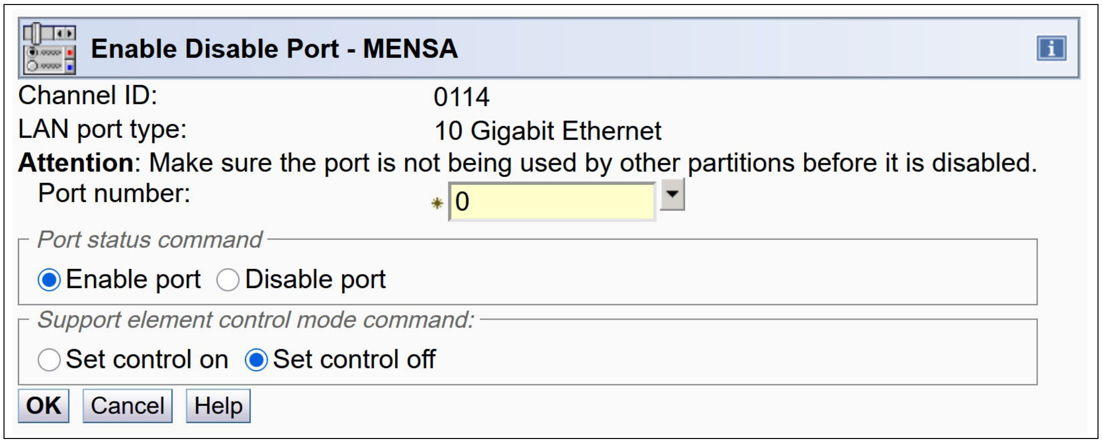

Figure 6-5 Enable or Disable Ports

On this window, you first have to select the port that you want to enable or disable. The port selection is shown always, even if the selected PCHID only has one port, like for the new Network Express Adapter. In such a case, the port selection will only offer the value “0”.

After selecting the port, you can decide if this port is to be enabled or disabled. And in addition you can decide to set the SE control to “on” or “off”. By setting this to “on”, the PCHID will only allow configuration commands to come from the service element.

Then click **OK** to execute your selection.

Note: On IBM z17 the option “Support element control mode command” is not available.

6.1.4 Set card mode

The configuration option “Set card mode” is only available for channel type OSD on OSA-Express7S 1000BaseT on IBM z15.

On IBM z16 and IBM z17, OSA Express7S 1000BaseT (FC 0446) only supports auto-negotiate. The OSA automatically negotiates the best speed (100mbps or 1000mbps) and duplex mode (Full or Half) with the switch/router that it’s connected to.

6.2 Verifying OSA configuration

For OSD as well as OSH channel types, you will find the option “Query port status” in the card specific advanced facilities as shown in Figure 6-6.

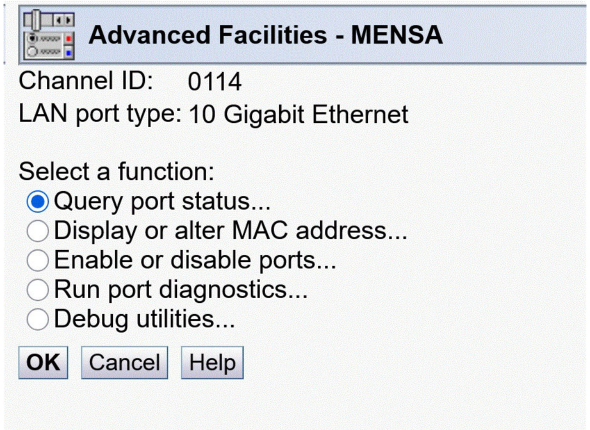

Figure 6-6 OSA Card specific advanced facilities

You can select this option and click OK to show the current port status as shown in Figure 6-7 on page 134.

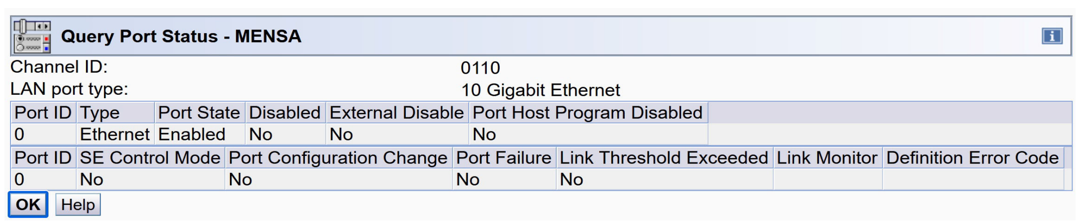

Figure 6-7 Query port status

6.3 Configuring Shared Memory Communications (SMC-D and SMC-R)

The concept of Shared Memory Communications (SMC) is explained in more detail in "Shared Memory Communication (SMC)" on page 35. SMC can be done either to target to a remote machine, which is called SMC-R using the RoCE (Remote Direct Memory Access (RDMA) over Converged Ethernet) functionality. Or you can target an LPAR within the same CPC with SMC using Internal Shared Memory (ISM). This is called SMC-D (Direct).

In both cases (SMC-R and SMC-D) the SMC connection has to be established using a TCP connection. Therefor, each SMC function has to be paired with a corresponding OSD, OSH[1], or IQD channel for the initiating TCP connection.

If you define a PNETID to the OSD or OSH channel used for establishing the TCP connection, the same PNETID has to be specified for the corresponding NETH or ISM function.

[1] For SMC-R paired with OSH, both definitions (OSH and NETH) have to be for the same PCHID.

6.3.1 Defining Shared Memory Communications over RoCE (SMC-R)

The SMC-R function for IBM z17 is provided by the Network Express Adapter only. The name of the RoCE PCIe function to be defined is NETH, and the definition of this function is explained in detail in 15.2.3, “Defining a NETH PCIe function” on page 339.

6.3.2 Defining Shared Memory Communications Direct (SMC-D)

The SMC-D function is provided by defining the Internal Shared Memory (ISM) function. It is described in detail in 15.2.2, “Defining an ISM PCIe function” on page 334.

7

Defining console communication

Naming: The IBM z17 system that are targeted by this publication consist of IBM z17 ME1(9175). Throughout this chapter, we might refer to this machine as IBM z17.

Note: Not all the following configuration items are necessarily required for your installation. In addition, the examples that are presented are not exhaustive.

This chapter includes the following topics:

- Preparing a console definition
- Defining the OSA-ICC
- Configuring a new OSA-ICC by using OSA Advanced Facilities
- Verifying the OSA-ICC definition

7.1 Preparing a console definition

A non-Systems Network Architecture (SNA) console is required for IPL, z/OS system operation, and management. Define at least one non-SNA console to your z/OS system. To define the non-SNA console, use the Open Systems Adapter-Express (OSA-Express) Integrated Console Controller (ICC) (OSA-ICC) function. For more information, see *Open Systems Adapter Integrated Console Controller User's Guide*, SC27-9003.

Note: We created our definitions by using the IBM® Personal Communications. You can use other emulation compatible products currently in the marketplace.

The following steps are required:

1. Define the OSC channel path ID (CHPID) and the CNTLUNIT.
2. Define the 3270-X IODEVICE.
3. Configure OSA-ICC by using the Open Systems Adapter (OSA) Advanced Facilities.
4. Export / import the OSA-ICC configuration (optional).
5. Activate the OSA-ICC configuration by using the OSA Advanced Facilities.

6. Set up IBM® Personal Communications.

> **Note:** Channel type OSC is supported on the OSA Express7S 1.2 GbE SX/LX adapters only.

Configure OSA-ICC when you upgrade to IBM z17 from an older IBM Z generation where OSA-ICC definitions exist or when installing a new IBM z17 and non-SNA consoles are required.

7.2 Defining the OSA-ICC

Before you perform the OSA-ICC configuration, define the OSC CHPID, CNTLUNIT, and 3270-X IODEVICE to I/O configuration using HCD. This definition is described in Chapter 13, “Adding network devices” on page 265.

7.3 Configuring a new OSA-ICC configuration by using OSA Advanced Facilities

When installing a new IBM z17, configure OSA-ICC from scratch. To create an OSA-ICC configuration, complete the following steps:

1. Log on to the Hardware Management Console (HMC), open **OSA Advanced Facilities** and select the central processor complex (CPC).
2. Select the OSC physical channel ID (PCHID) to use for the OSA-ICC configuration, and select **Card Specific Advanced Facilities** (Figure 7-1).

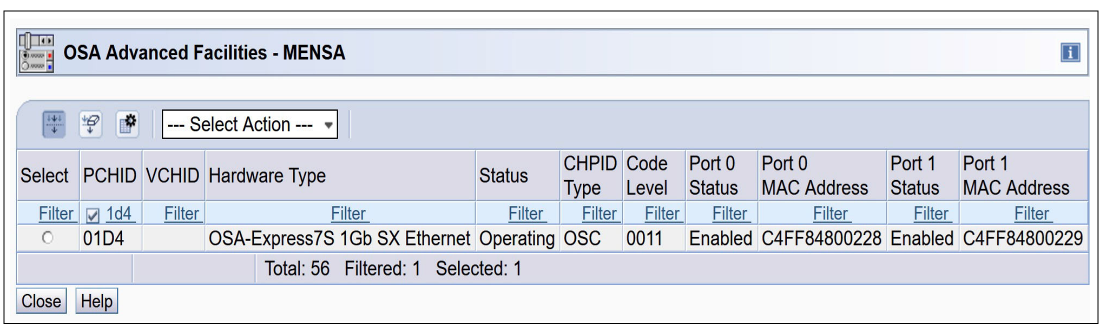

Figure 7-1 HMC: Card Specific Advanced Facilities

> **Note:** From z17, there is a new column in OSA Advanced Facilities called VCHID. OSC channel does not exploit VCHID definition.

3. Select **Panel configuration options** and click **OK** (Figure 7-2 on page 139).

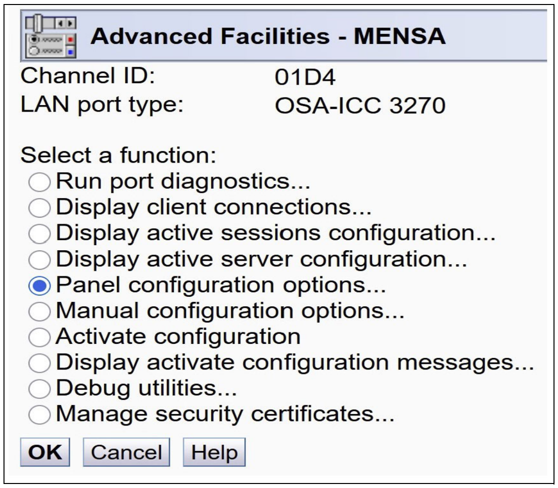

Figure 7-2 HMC: Panel configuration options

4. The Panel Configuration Options window opens (Figure 7-3). Define the session and server configurations in this window, and validate those values. In this example, we define the server configuration first, and then define the session configuration. To edit the server configuration, select **Edit server configuration** and click **OK**.

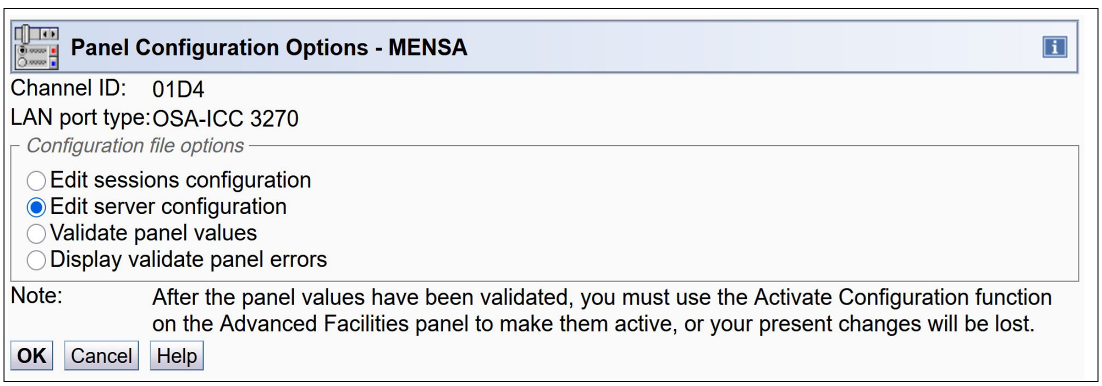

Figure 7-3 HMC: Edit server configuration

5. The Edit Server Configuration window opens. Enter the necessary values on this window. Figure 7-4 shows our sample configuration. Click **OK** to save.

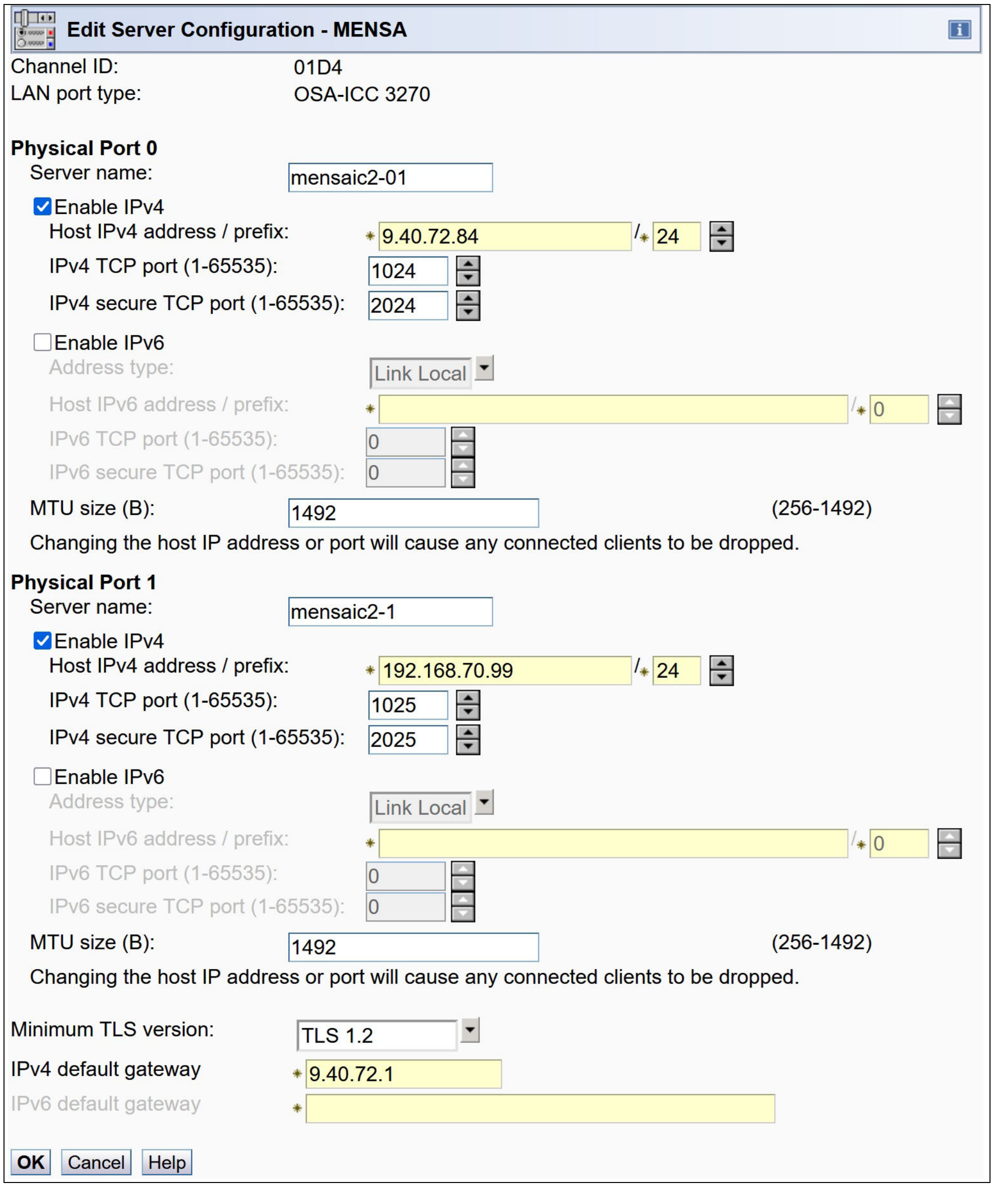

Figure 7-4 HMC: Edit Server Configuration window

> **Note:** Starting with IBM z17, OSC can exploit TLS 1.3 support.

For more information, see the *OSA-Express Integrated Console Controller Implementation Guide*, SG24-6364.

6. The command is completed (ACT20402 and the window opens. Click **OK** to return to the Panel Configuration Options window.

7. Next, define the session configuration. Select **Edit sessions configuration** and click **OK**. The window that is shown in Figure 7-5 opens.

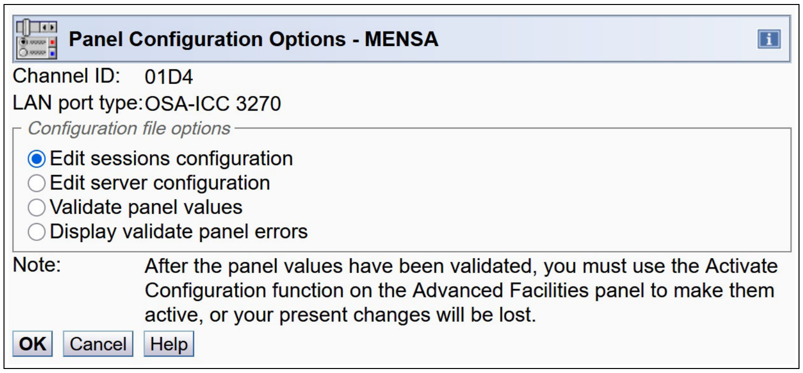

Figure 7-5 Edit sessions configuration

8. The Edit Sessions Configuration window opens (Figure 7-6). To configure a session, select a number from the Session Index column and click **Change**.

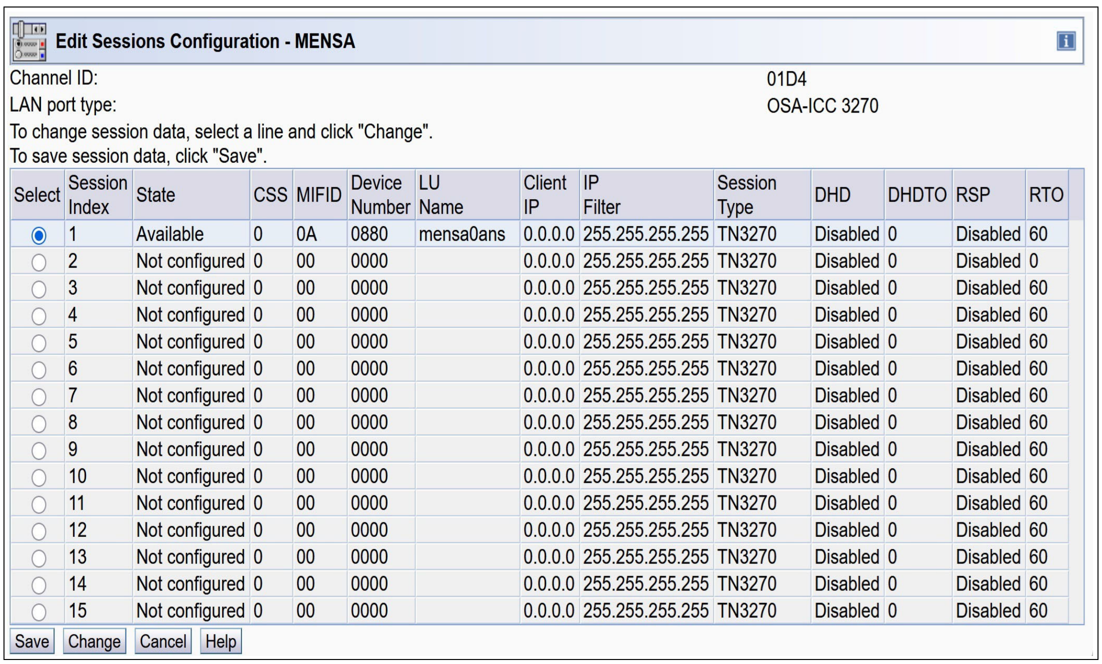

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
◉	1	Available	0	0A	0880	mensa0ans	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	2	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	0
○	3	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	4	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	5	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	6	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	7	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	8	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	9	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	10	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	11	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	12	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	13	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	14	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	15	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Figure 7-6 HMC: Edit Sessions Configuration selection

9. The Edit Session Configuration window opens. Define the session parameter here. Click **OK** to save (Figure 7-7).

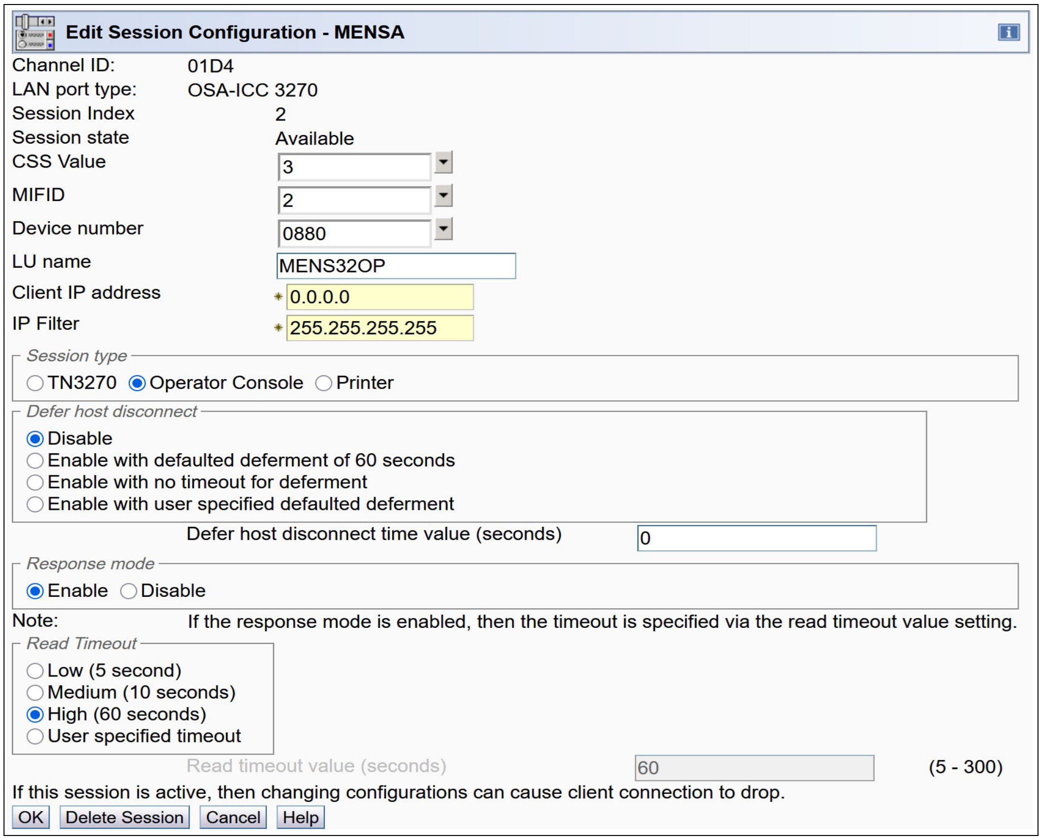

Figure 7-7 HMC: Edit Session Configuration input

10. The Edit Sessions Configuration window opens again (Figure 7-8). Be sure that your input values are displayed correctly. To save the session values, click **Save**.

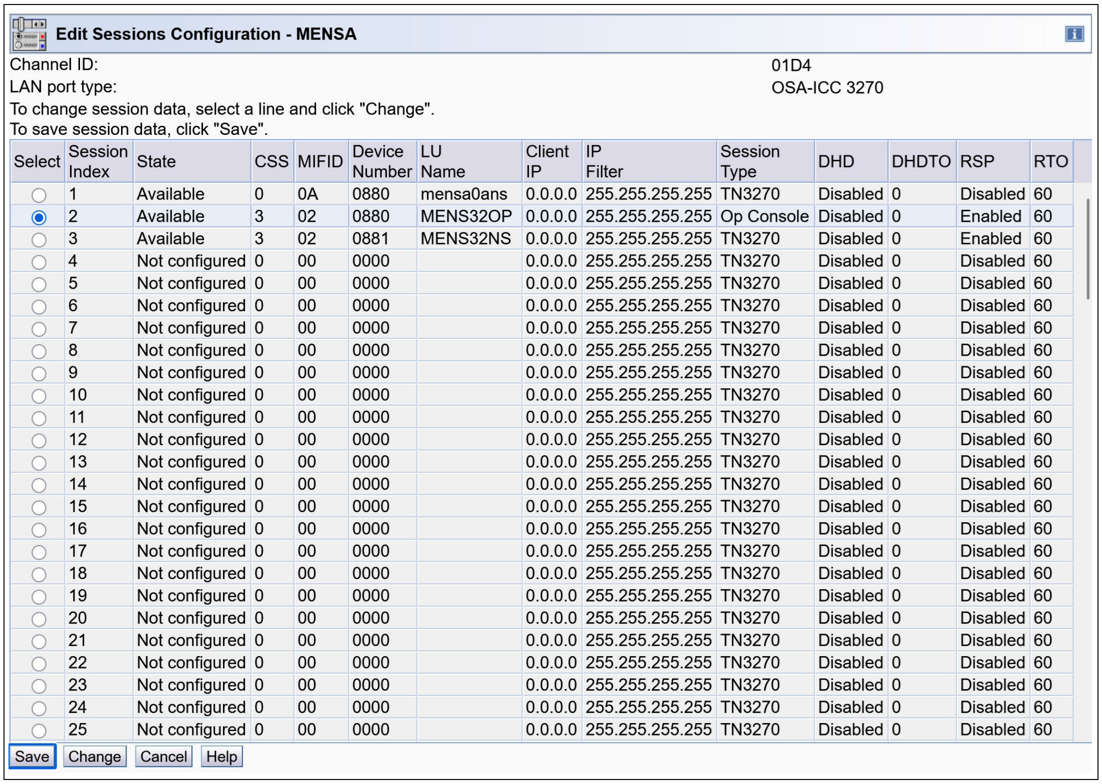

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
○	1	Available	0	0A	0880	mensa0ans	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
●	2	Available	3	02	0880	MENS32OP	0.0.0.0	255.255.255.255	Op Console	Disabled	0	Enabled	60
○	3	Available	3	02	0881	MENS32NS	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	4	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	5	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	6	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	7	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	8	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	9	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	10	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	11	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	12	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	13	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	14	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	15	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	16	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	17	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	18	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	19	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	20	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	21	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	22	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	23	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	24	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	25	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Figure 7-8 HMC: Edit Sessions Configuration after you define the values

11. The command completed (ACT20402) window opens. Click **OK**.

12. The Panel Configuration Options window opens again (Figure 7-9). Now, you can validate the values. Select **Validate panel values** and click **OK**.

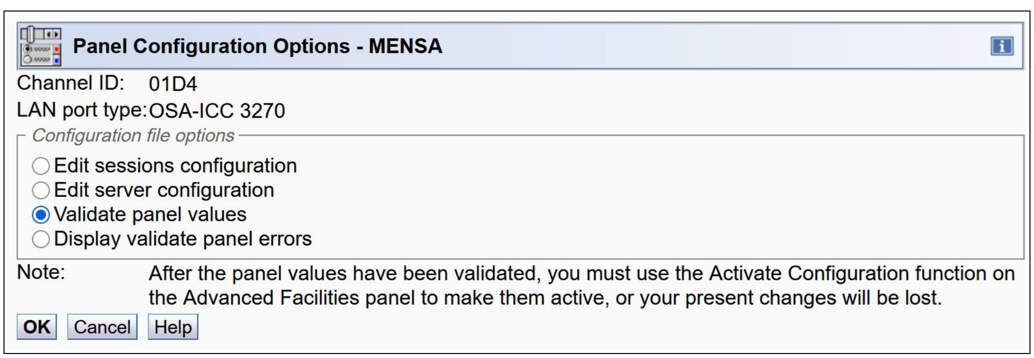

Figure 7-9 HMC: Validate panel values

13. If the configuration data is correct, the command completed (ACT20402) window opens. If an error is encountered, a window like Figure 7-10 on page 144 opens. You can confirm the error by selecting **Display validate panel errors**, as shown in Figure 7-11 on page 144, in Panel Configuration Options (Figure 7-12 on page 144). Correct the error, and select again **Validate panel values** to recheck.

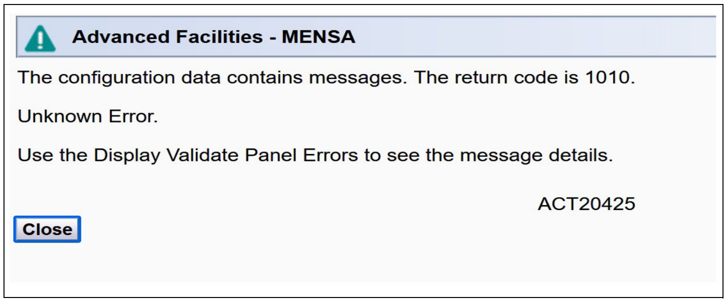

Figure 7-10 HMC: Panel Configuration Options error found

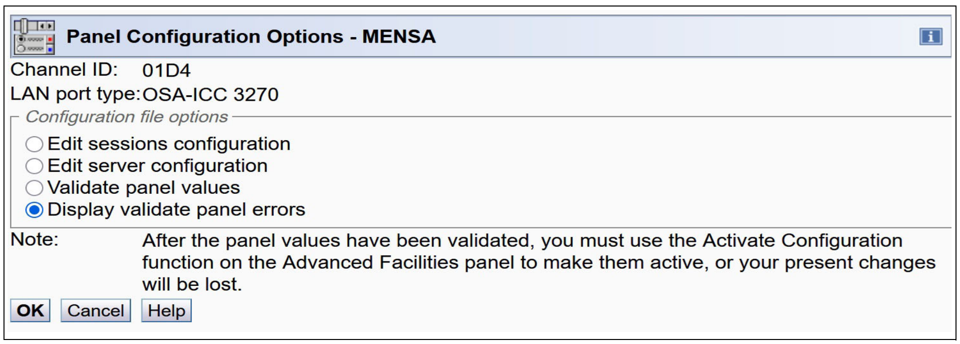

Figure 7-11 HMC: Display validate panel errors

Figure 7-12 shows the details of the error message.

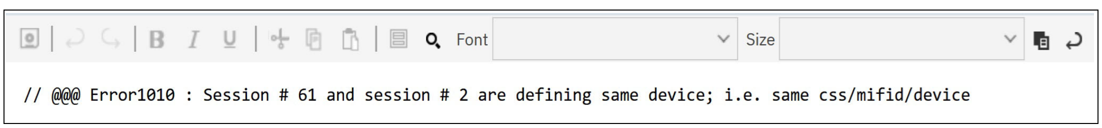

Figure 7-12 HMC: Detailed error message.

14. When the validation is complete without any errors, you can activate the OSA-ICC configuration. To activate the configuration, exit the Panel Configuration Options window by clicking **Cancel**, and then select **Activate configuration** in the Advanced Facilities window (Figure 7-13).

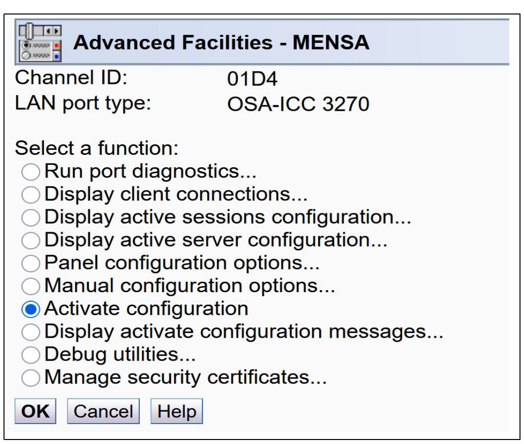

Figure 7-13 HMC: Activate configuration

15. The confirmation window opens. Click **Yes** to continue (Figure 7-14).

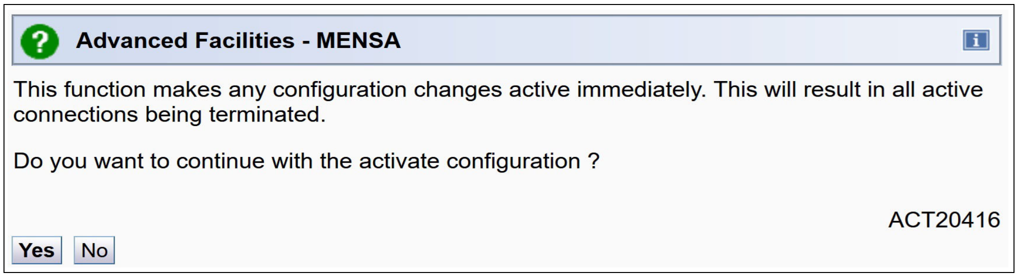

Figure 7-14 HMC: Activate configuration confirmation

16. When the activation is complete, the message that is shown in Figure 7-15 appears.

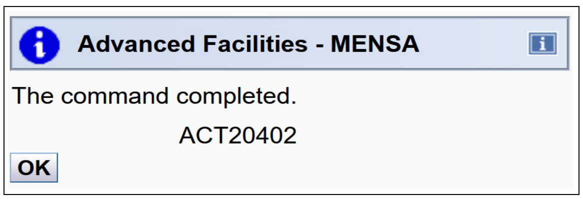

Figure 7-15 HMC: Activate configuration success

7.3.1 Saving and restoring the OSA-ICC configuration

During an upgrade, you can export an OSA-ICC configuration file from the source IBM Z CPC and import it to the IBM z17. This section describes how to export and import the OSA-ICC configuration file by using HMC and an FTP server that is part of the same local area network (LAN) segment.

Exporting the OSA-ICC configuration file by using OSA Advanced Facilities

In this example, we export the OSA-ICC configuration file from an IBM z17 to a FTP server and import the file to the same IBM z17 from the same FTP server. You can also import or export the configuration file by using a USB Flash Drive.

To export the OSA-ICC configuration file, complete the following steps:

1. Before you try to export a source file to an FTP server, make sure the FTP server is reachable from this particular HMC:
   a. Contact your LAN administration and ask for the hostname and TCP/IP address of the FTP server that is connected to the same subnet. From the HMC Welcome window, click **HMC Management**, as shown in Figure 7-16 on page 146.

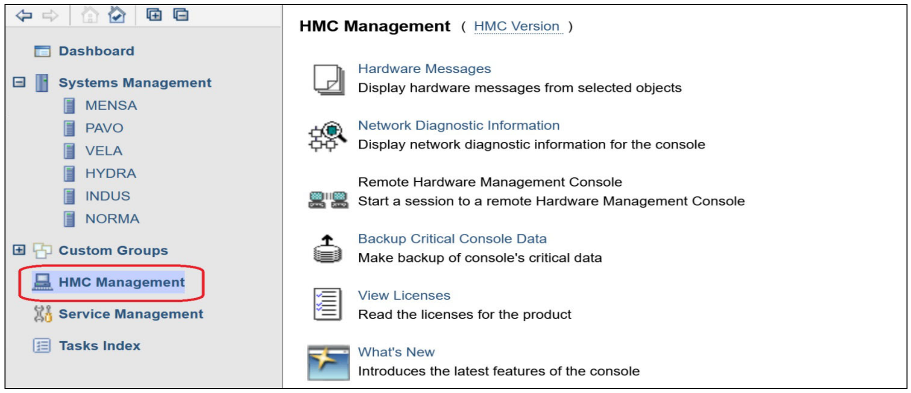

Figure 7-16 HMC: Clicking HMC Management to get to the network diagnostic information

b. On the HMC Management window, click **Network Diagnostic Information** to get to the Network Diagnostic Information window, as shown in Figure 7-17.

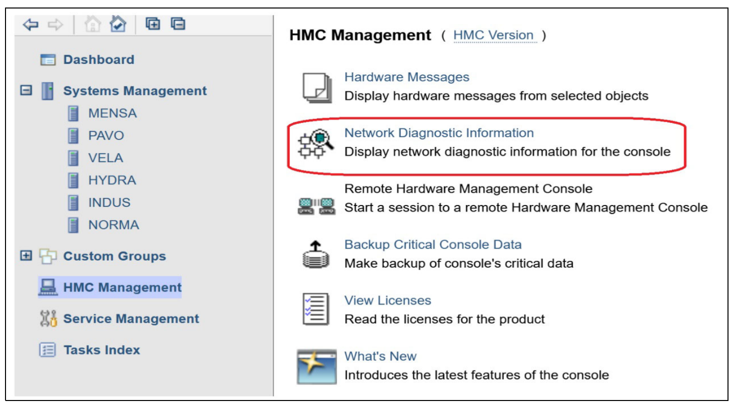

Figure 7-17 HMC: Clicking Network Diagnostic Information to ping the FTP server

c. On the **Ping** tab of the Network Diagnostic Information window, enter the TCP/IP Address of the FTP server and click **Ping**. You should see the ping statistics, which indicate that this HMC has a connection to the FTP server, as shown in Figure 7-18.

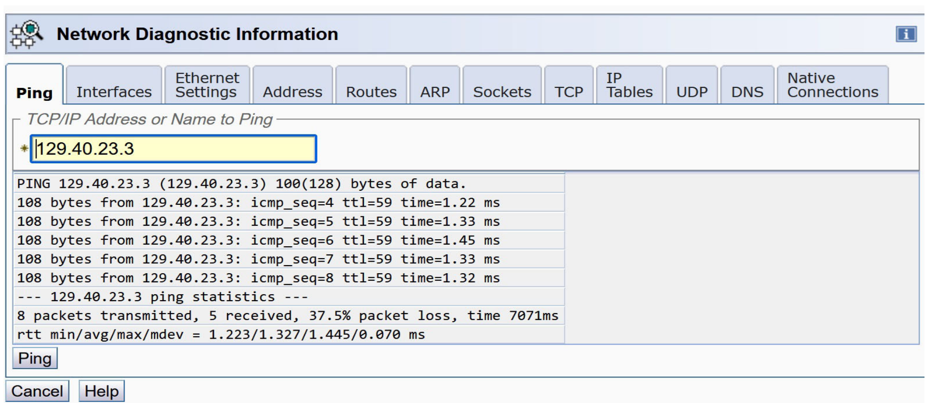

Figure 7-18 HMC - Network Diagnostic Information: Using the ping function to verify the physical connection between this HMC and the FTP server

2. Log on to the HMC, select the CPC that you want to operate, and open the OSA Advanced Facility.
3. Select the OSC CHPID to export the OSA-ICC configuration file, select **Card Specific Advanced Facilities**, and then select **Manual configuration options**. Click **OK** (Figure 7-19).

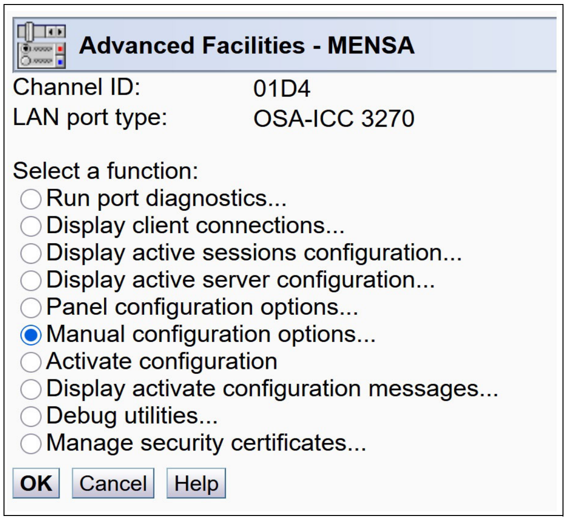

Figure 7-19 HMC: Manual configuration options

4. The Manual Configuration Options window opens (Figure 7-20). Select **Export source file by FTP** and click **OK**.

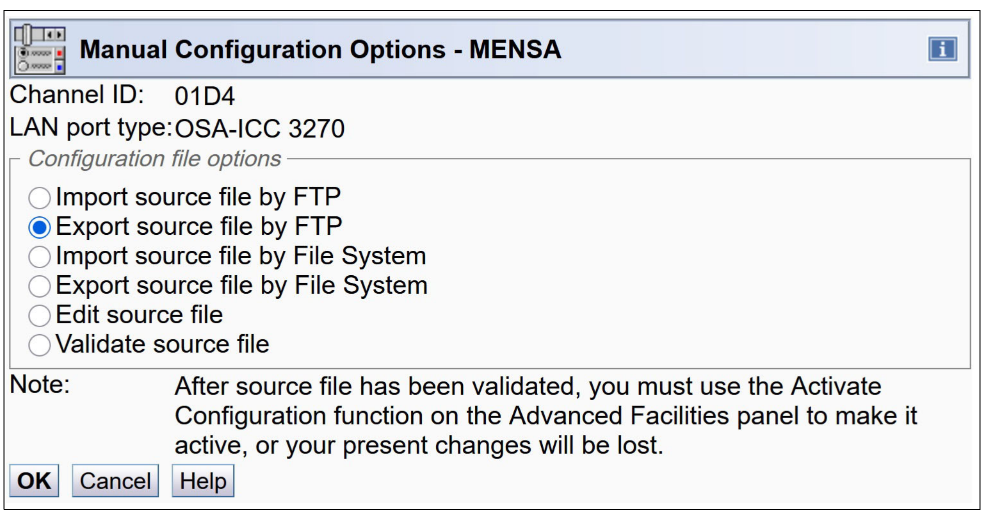

Figure 7-20 HMC: Export source file by FTP

5. The window prompts you for the FTP server information and the location of the file to export. For our example, we enter `osaicc_1D4.txt`, as shown in Figure 7-21 on page 148. Click **Export**. The HMC task writes the source file for the PCHID that was selected on to the FTP server and displays a message when it completes (Figure 7-22 on page 148). Click **OK**.

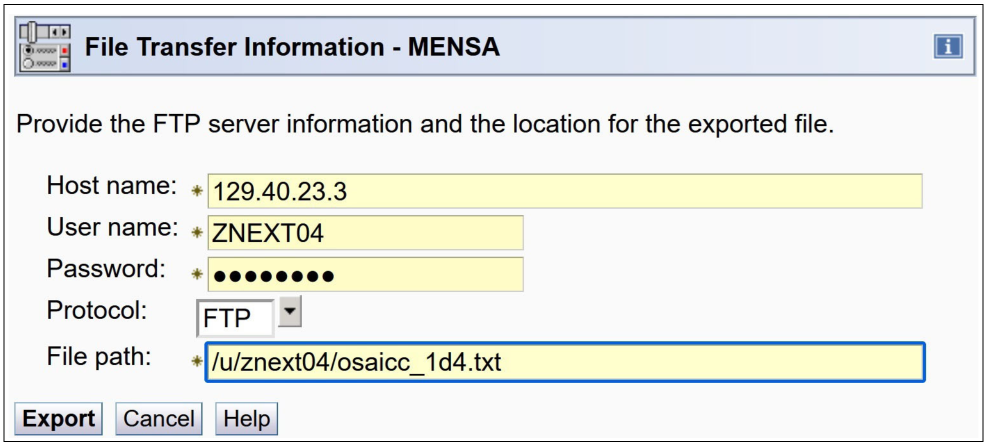

Figure 7-21 HMC - Export file: Specifying the FTP server, and file path and name

The HMC displays the ACT20402 window. Click **OK**.

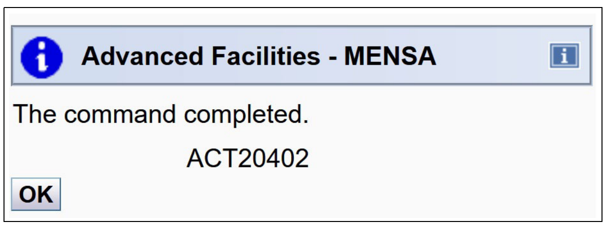

Figure 7-22 HMC: ACT20421 window display

6. Click **Cancel** to exit all OSA Advanced Facilities windows.

Example 7-1 shows an extract from the source file that we transferred to the FTP server.

Example 7-1 OSA-ICC: Sample source file

```
<OSC_SERVER>
<OSC_PHYSICAL_PORT0>
  HOST_IP= 9.40.72.84
  SUBNET_MASK= 255.255.255.0
  PORT= 1024
  SECURE_PORT= 2024
  HOST_LL_ADDRESS/PREFIX=  fe80::c6ff:84ff:fe80:228/64
  ADDR_TYPE= LINK_LOCAL
  HOST_IPV6_ADDRESS/PREFIX=  ::/0
  IPV6_PORT= 0
  IPV6_SECURE_PORT= 0
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= mensaic2-01
</OSC_PHYSICAL_PORT0>

<OSC_PHYSICAL_PORT1>
  HOST_IP= 192.168.70.99
  SUBNET_MASK= 255.255.255.0
  PORT= 1025
  SECURE_PORT= 2025
  HOST_LL_ADDRESS/PREFIX=  fe80::c6ff:84ff:fe80:229/64
```

```
  ADDR_TYPE= LINK_LOCAL
  HOST_IPV6_ADDRESS/PREFIX=  ::/0
  IPV6_PORT= 0
  IPV6_SECURE_PORT= 0
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= mensaic2-1
</OSC_PHYSICAL_PORT1>

  TLS_VERSION= 1.2
  DEFAULT_GATEWAY= 9.40.72.1
  IPV6_DEFAULT_GATEWAY=  ::
</OSC_SERVER>

<CONFIG_SESSION>
<SESSION1>
  CSS= 00 IID= 0A DEVICE= 0880
  GROUP=  "mensa0ans"
  CONSOLE_TYPE= 1    RESPONSE= OFF    READ_TIMEOUT= 60
</SESSION1>
<SESSION2>
...
...
</CONFIG_SESSION>
```

Editing the source file for OSA-ICC

When your OSA-ICC configuration for IBM z17 (such as the IODEVICE, CSSID, and MIFID of logical partitions (LPARs)) is changed, you must edit the OSA-ICC source file to match the new configuration. For more information about editing the source file, see *Open Systems Adapter Integrated Console Controller User's Guide,* SC27-9003.

Importing the OSA-ICC source file to IBM z17

To import the source file and activate the configuration of the OSA-ICC on the new IBM z17, complete the following steps:

1. Before you import the source file from the FTP server, make sure that the FTP server is reachable from this particular HMC. For more information about how to make sure that the FTP server is reachable, see step 1 on page 145.
2. Log on to the HMC, select the CPC that you want to operate, and open OSA Advanced Facility.
3. Select OSC CHPID to import the OSA-ICC configuration file. Then, select **Card specific advanced facilities**.
4. Select **Manual configuration options** and click **OK**.
5. The Manual Configuration Options window opens. Select **Import source file by FTP** and click **OK** (Figure 7-23 on page 150).

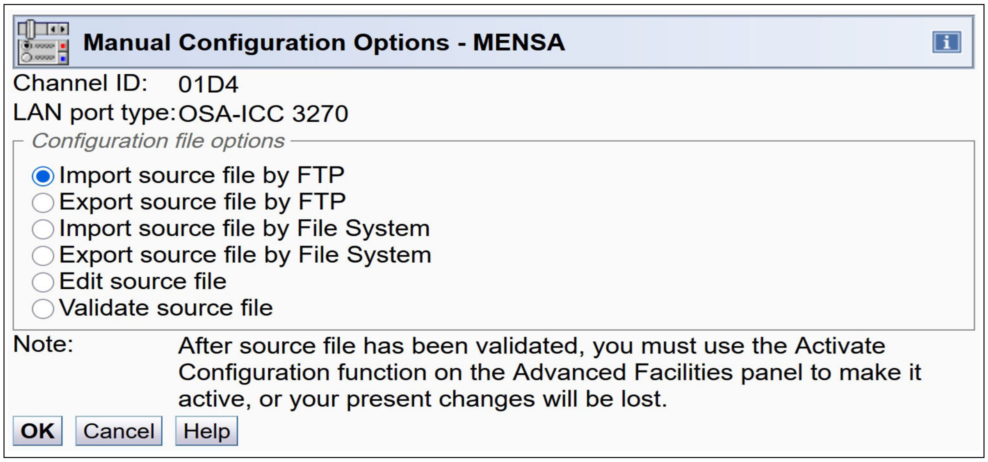

Figure 7-23 HMC: Importing a source file

6. You are prompted to provide the FTP server and file location information (Figure 7-24). Click **Import**.

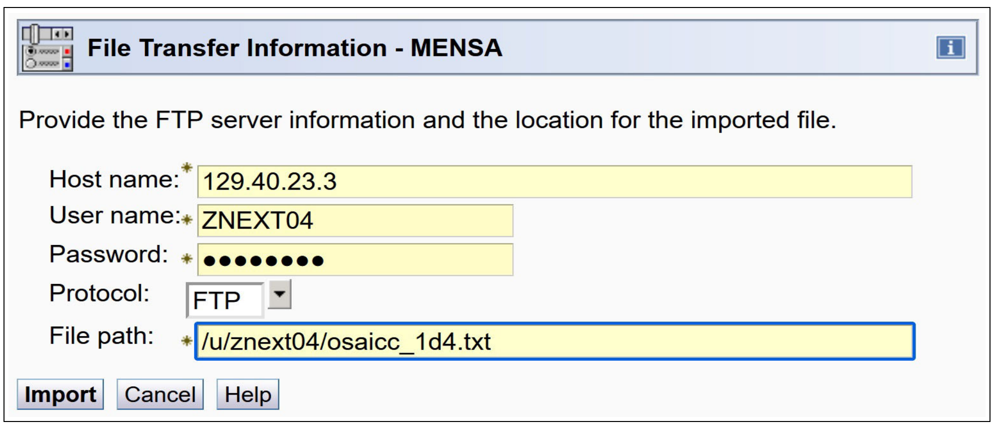

Figure 7-24 HMC: Importing a file

7. Figure 7-25 indicates that the source file import is complete. Click **OK** to continue.

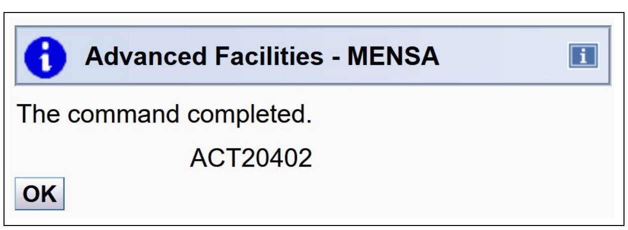

Figure 7-25 HMC: Importing file successful

8. After importing the source file, you must validate it and activate the configuration.

7.4 Verifying the OSA-ICC definition

This section presents some commands that you can use to verify the OSA-ICC configuration. Before starting, the system programmer must activate the new I/O to make changes to a running configuration.

7.4.1 z/OS commands

You can check your definition by using the following z/OS commands:

- **DISPLAY M=CHP(xx)**

  Check whether the `CHPID DESC` is displayed as `OSA CONSOLE` (Figure 7-26).

```
D M=CHP(B1)
IEE174I 21.21.33 DISPLAY M 767
CHPID B1:  TYPE=14, DESC=OSA CONSOLE, ONLINE
DEVICE STATUS FOR CHANNEL PATH B1
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
0088 +  +  $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@ $@
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 01D4
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 7-26 OSC D M=CHP

- **DISPLAY M=DEV(xxxx)**

  Using this command, you can confirm the channel path to a device. The node descriptor information that is returned includes the emulated control unit (CU) 1B10. Also included is the 9175 machine type and serial number, as shown in Figure 7-27.

```
D M=DEV(880)
IEE174I 21.22.06 DISPLAY M 769
DEVICE 00880   STATUS=ONLINE
CHP                   B1
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             1B10
INTERFACE ID          D400
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = NOT AVAILABLE
SCP TOKEN NED     = 002074.   .IBM.02.9175000B9FB8.B100
SCP DEVICE NED    = 002074.002.IBM.02.9175000B9FB8.B100
```

Figure 7-27 OSC D M=DEV(xxx)

- **DISPLAY U**

  Using this command, you can confirm the device number and the status. Check whether the device number and type are defined correctly (Figure 7-28).

```
D U,,,880,1
IEE457I 21.22.44 UNIT STATUS 771
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
0880 3270 O                                       0
```

Figure 7-28 OSC D U,,,device

7.4.2 OSA-ICC console initial window

When you complete setting up IBM Personal Communications, you see the OSA-ICC initial window, as shown in Figure 7-29. Check whether the OSA-ICC definitions are reflected correctly.

If you do not see the initial window, check your definition in the OSA-ICC or IBM Personal Communications session setup.

```
 **    OSC Index 02 connected to mensaic2-01 via IP Addr 9.40.72.84:1024     **
 **        LT Index=06 CSSID=03 MIFID=02 CU=0  UA=00  LUName=MENS320P        **
 **       Type=9175-ME1 Mfg=IBM SN=0000000B9FB8 CHPID=B1 Status=Active       **



9.67.186.61:63237
```

Figure 7-29 OSA-ICC initial window

8

Preparing for IBM Parallel Sysplex and Server Time Protocol

This chapter describes the preparation tasks for IBM Parallel Sysplex and how to configure time synchronization with Server Time Protocol (STP).

Note: The examples that are shown in this chapter are based on the IBM z17 ME1 (9175).

This chapter includes the following topics:

- Preparing for Parallel Sysplex
- Preparing for non-sysplex system time synchronization
- Server Time Protocol overview
- Configuring the HMC as an NTP server
- HMC 2.17.0 (Manage System Time task)

8.1 Preparing for Parallel Sysplex

If your IBM z17 is a member of a Parallel Sysplex or if you create a Parallel Sysplex that involves an IBM z17, time synchronization among central processor complexes (CPCs) is required. For time synchronization, use the STP feature (Feature Code 1021).

In addition to the time synchronization, there are many steps to be followed in order to create a Parallel Sysplex. You can find more details in the manual *MVS Setting Up a Sysplex*, SA23-1399-60.

Some of those steps are performed when you are installing an IBM z17:

1. Define Coupling Facility (CF) links.
2. Define Fibre Connection (FICON) channel-to-channel (FCTC) connections (optional).
3. Define the CF logical partition (LPAR).
4. Define the CF LPAR image profile in Hardware Management Console (HMC).
5. Configure STP.

You can find more details on this redbook about the items 1 to 4 above on the following chapters:

- Defining CF (coupling) links is described in Chapter 9, “Defining Coupling Facility links” on page 165.
- Defining CTC is described in Chapter 16, “Defining Channel to Channel Connections” on page 353.
- Defining the CF logical partition (LPAR) is described on 14.2.1, “Defining Coupling Facility LPARs in a channel subsystem” on page 304
- Defining the image profile is described in 5.4, “Creating an image profile on the Support Element” on page 105.

8.2 Preparing for non-sysplex system time synchronization

If your IBM z17 is not part of a Parallel Sysplex but you want to synchronize the time among multiple CPCs, you need the STP function and timing-only links. These configuration steps are required:

1. Define timing-only links.
2. Configure STP.

Defining timing-only links is described in 9.4, “Defining an STP timing-only link by using ICA SR” on page 183.

8.3 Server Time Protocol overview

STP provides time synchronization among multiple CPCs in a Coordinated Timing Network (CTN). A CTN is a collection of servers that are synchronized to a time protocol that is called Coordinated Server Time (CST). The STP function (feature) is implemented in the Licensed Internal Code (LIC) as a chargeable feature. You must order STP enablement Feature Code 1021 to use STP.

In a timing network that is based on STP, we have the concept of stratum levels to define the hierarchy of a server in the timing network:

- A stratum 1 server is the highest level in the hierarchy of timing network that uses STP to synchronize to the Coordinated Server Time (CST). The stratum 1 server has a direct connection to the PTP/NTP server(s).
- Stratum 2 and stratum 3 levels are determined by how many stratum levels they are away from stratum 1.
- A server that uses STP messages to synchronize to a stratum 1 server is referred to as a stratum 2 server.
- Similarly, a server that uses STP messages to synchronize to a stratum 2 server is referred to as a stratum 3 server.
- Stratum level 4 servers uses STP messages to synchronize to a stratum 3 server and have no direct STP link connections to the stratum level 1 server (Current Time Server - CTS).

For more information about STP concepts and planning information, see *IBM Z Server Time Protocol Guide*, SG24-8480.

Beginning with IBM z15, STP stratum level 4 is supported. Timekeeping information is transmitted over coupling links.

Note: Stratum 4 is a temporary status to allow more options to move and replace machines and reconfigure CTN. Although STP stratum level 4 is supported, it should not be used for permanent configurations. Stratum 4 should be used for transitional configurations during CTN maintenance.

Figure 8-1 shows a diagram of a CTN.

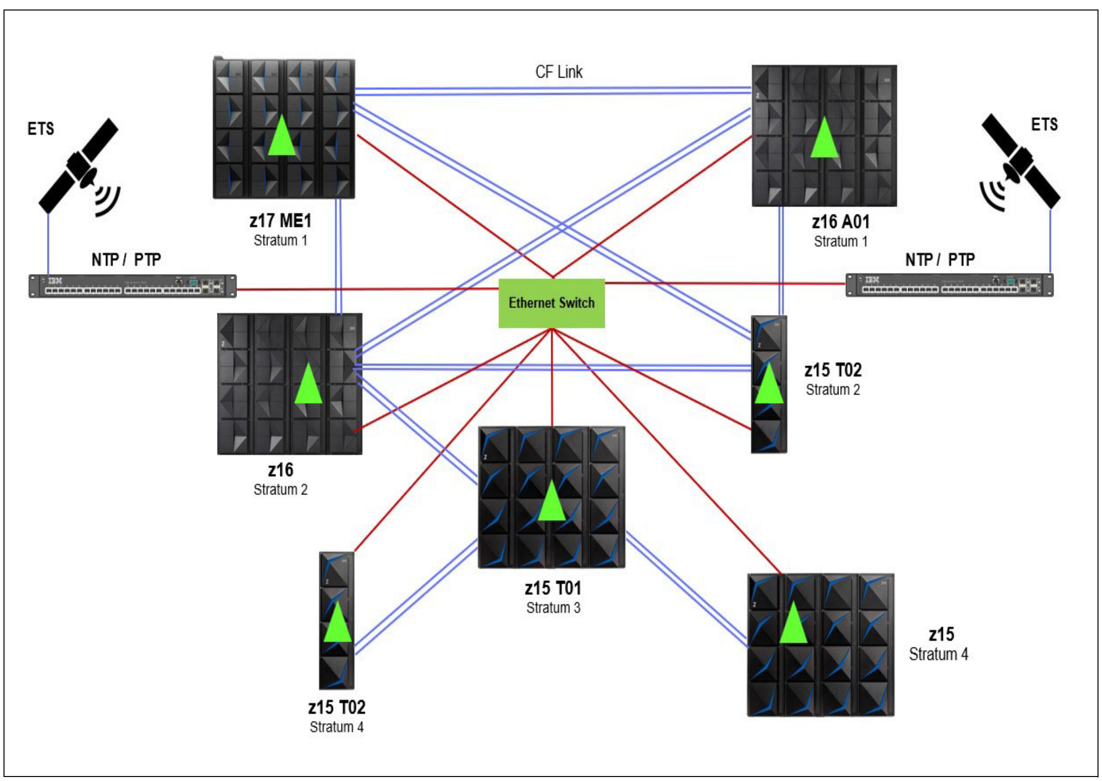

Figure 8-1 Diagram of a CTN

Note 2: On the CTN diagram presented previously, it is showed only one Ethernet Switch just to keep the figure simpler. It is recommended to have at least two on your configuration for availability reasons.

8.3.1 External Time Source

For IBM Z machines before IBM z15, the external time reference would typically be Network Time Protocol (NTP). Starting with IBM z15, support for IEEE 1588 (PTP) was added. Before IBM z16, the external time reference network connectivity occurred through the internal Support Element (SE) Management network Interface.

With IBM z17 and z16 you can access the External Time Source (ETS) by connecting the CPC directly to the client network, which provides access to the PTP or NTP time reference, thus bypassing the SE and the internal network. PTP or NTP traffic is provided to the CPC through Ethernet connectivity to the IBM z17 CPC drawer. The hardware features timestamping capabilities so that IBM z17 leverages the accuracy of PTP. In IBM z17, a container firmware partition provides time protocol support for both NTP and PTP.

For clients requiring a Coordinated Universal Time lower than 100 milliseconds, Pulse Per Second (PPS) connectivity is also available. To provide NTP or PTP data, the NTP or PTP servers are directly connected to the CPC.

Note: With IBM z17, a new BMC/OSC card has been implemented. There are two BMC/OSC cards per CPC drawer:

- Each CPC drawer has two combined BMC/OSC cards, each with one PPS port and one ETS port (RJ45 Ethernet, for both PTP and NTP).
- For PPS signal redundancy, two PTP/PPS ports must be connected. The point-to-point connections should not exceed 150m.

Figure 8-2 shows a sample PTP ETS configuration for the STP.

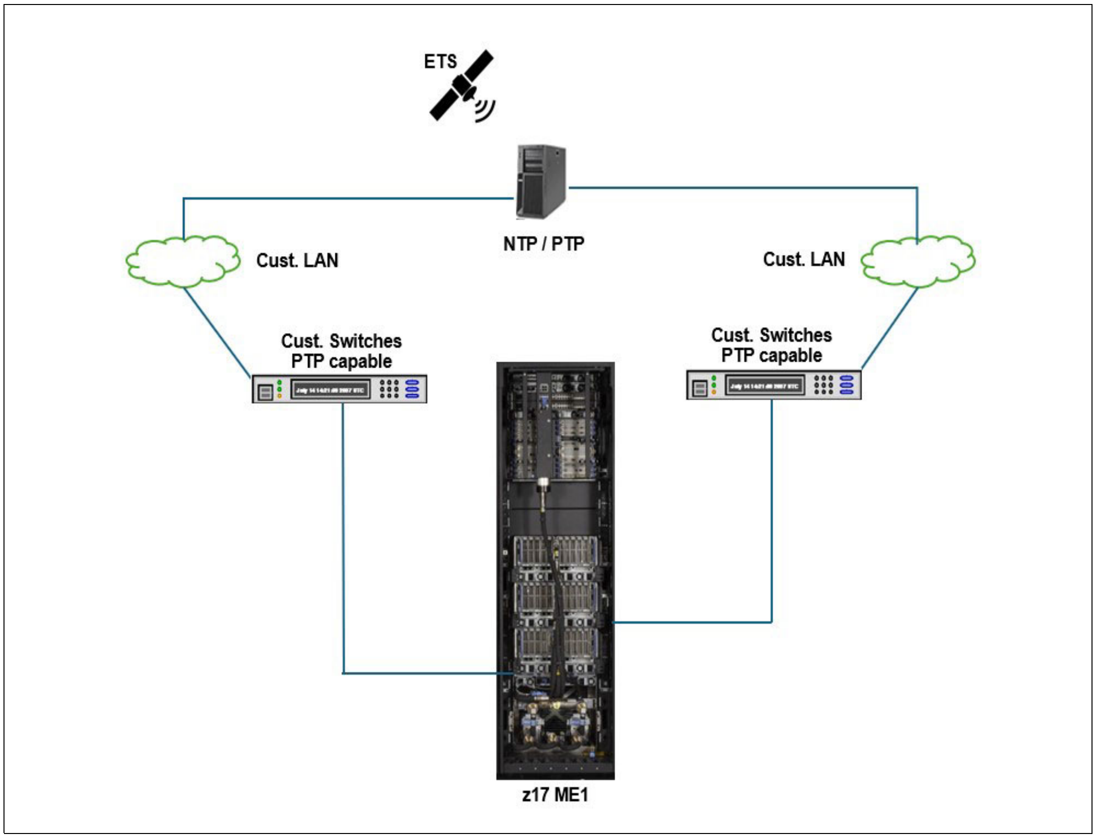

Figure 8-2 Example configuration: PTP time server configuration

For more information about this topic, see the white paper “STP PTP/NTP Time Direct to CEC Accuracy Performance”.

8.3.2 Configuring the HMC as an NTP server

On IBM z17, up to three NTP and two PTP servers can attach directly to the CPC. In addition, NTP can be secured with Network Time Security (NTS). This configuration would be considered as best practice, giving you the highest accuracy and availability for external time sources.

Still, there might be potential security concerns to directly connect the CPC to a network other than the often physically isolated SE LAN. For these cases, although not recommended, you can still define an NTP server on the HMC and use this as ETS for your CTN.

The HMC has two LAN ports that are physically isolated. One port is connected to the SE LAN while the other port can be used to connect to an external NTP server.

So, the NTP server on the HMC can access another NTP server through a separate LAN connection to obtain its time reference (see Figure 8-3). The NTP server function on the HMC does only provide NTP, not PTP or PPS. This configuration, leaves you with less accuracy and redundancy than the recommended setup with multiple directly connected time sources.

Note: As a best practice for security reasons, connect the ETS ports to a protected network (through a firewall) to access the external time servers (PTP or NTP). This best practice also applies for the ETS connectivity direct to the CPC.

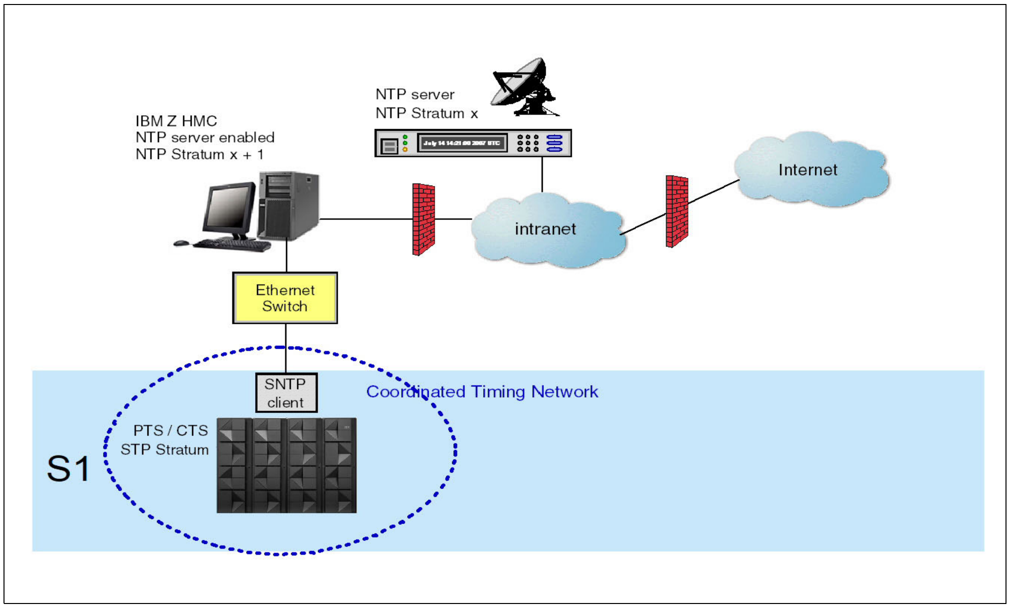

Figure 8-3 HMC configured as an NTP server

For more information, see *IBM Z Server Time Protocol Guide*, SG24-8480.

8.4 HMC 2.17.0 (Manage System Time task)

You can set up STP by using the HMC Manage System Time task. Figure 8-4 shows the Manage System Time task initial window on the HMC. The CTN configuration for any CPC object that is managed by this HMC can be displayed. You can display any CTN by clicking the down arrow icon next to the CTN ID name. You can enter the configuration wizard by using menu below the STP ACTIONS.

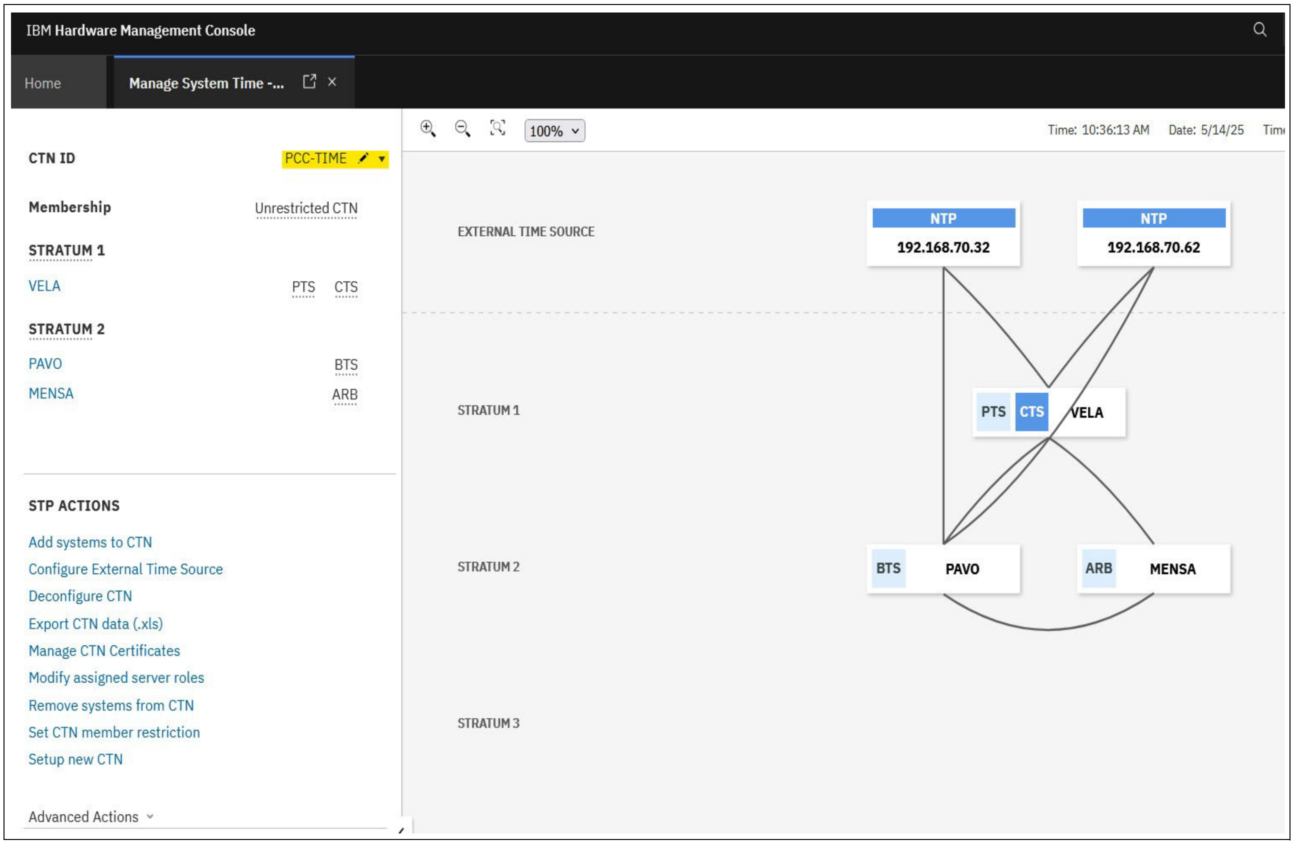

Figure 8-4 Display of a CTN with two external time sources

> **Note:** NTP servers configured as an ETS show their IP addresses instead of server names.

In the Manage System Time initial window, the CTN topology appears as a graph. You can identify the stratum level and the role of the CPC. You can also identify the status of each CTN component, such as CPC, CF LINK, and ETS, by clicking the respective objects like shown in Figure 8-5 on page 160 for a CPC and in Figure 8-6 on page 160 for a CF link.

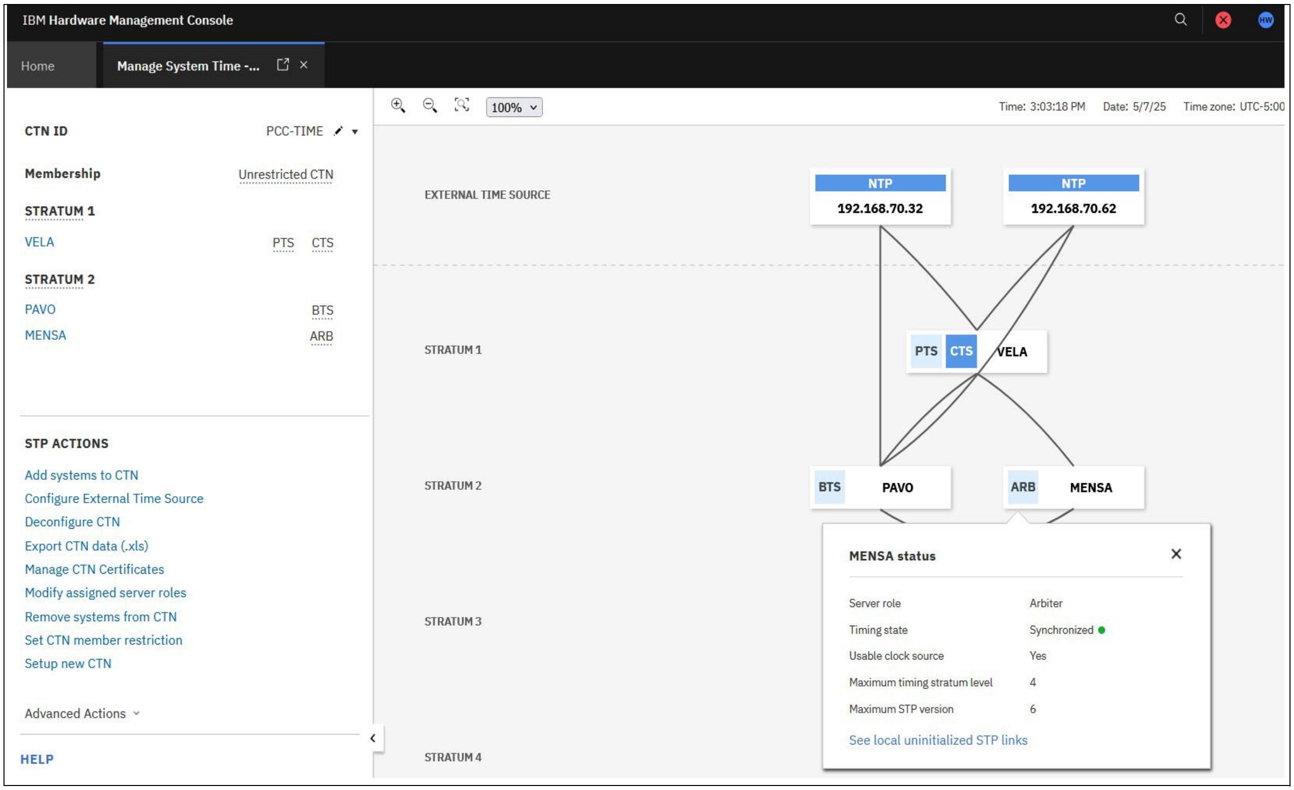

Figure 8-5 CPC status display

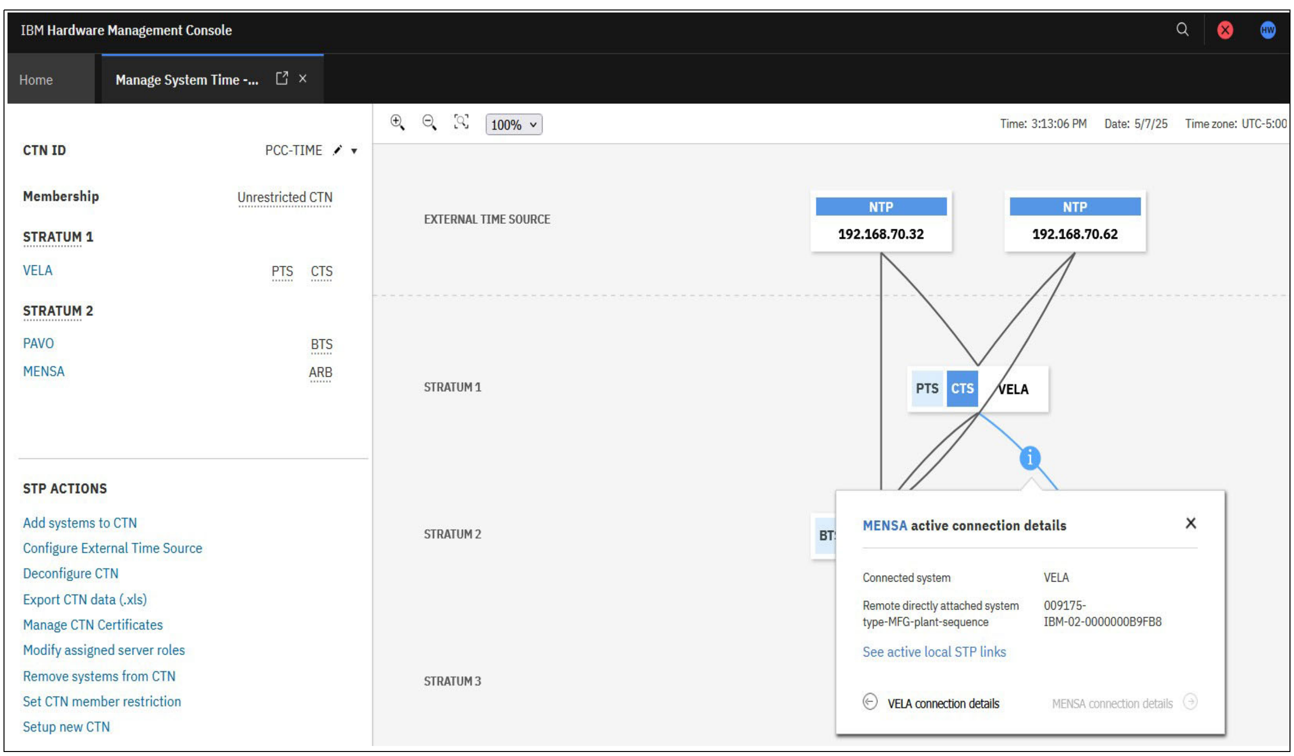

Figure 8-6 Coupling / timing link status display

8.4.1 Setting CTN member restrictions

For a single-server CTN or a dual-server CTN, you can restrict the CTN membership by selecting **Only allow the server(s) specified above to be in the CTN** in this menu. This configuration saves the configuration across Power on Resets (PORs) for STP-only CTNs with one or two servers (also known as *bounded CTN*).

8.4.2 HMC operations to add a CPC to the CTN

To add a CPC to an existing STP-only CTN, complete the following steps:

1. Open the **Manage System Time** task from the HMC, and select **Add CPC to CTN** from **STP ACTIONS** (Figure 8-7).

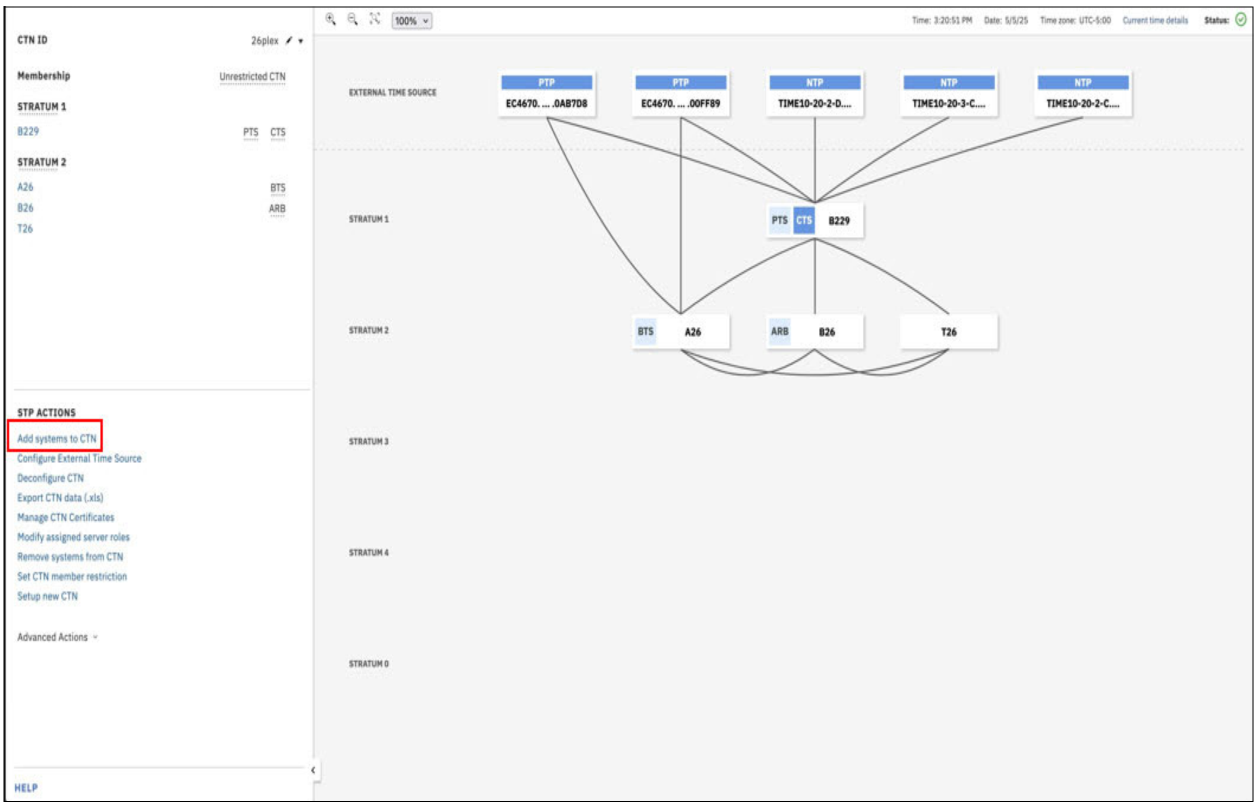

Figure 8-7 Add a CPC to a CTN

> **Note:** Before adding a CPC to an existing CTN, ensure that **Allow any server to be a member of the CTN** is selected in the **Coordinated Timing Network (CTN) member restriction preferences**. The “Membership” information should show “Unrestricted CTN”.

2. The Specify Coordinated Timing Network (CTN) Members window opens. In this window, select the CPC name that does not belong to any CTN and shall be added to your CTN. In this example, we chose CPC **A26** (Figure 8-8). Click “Next” to continue.

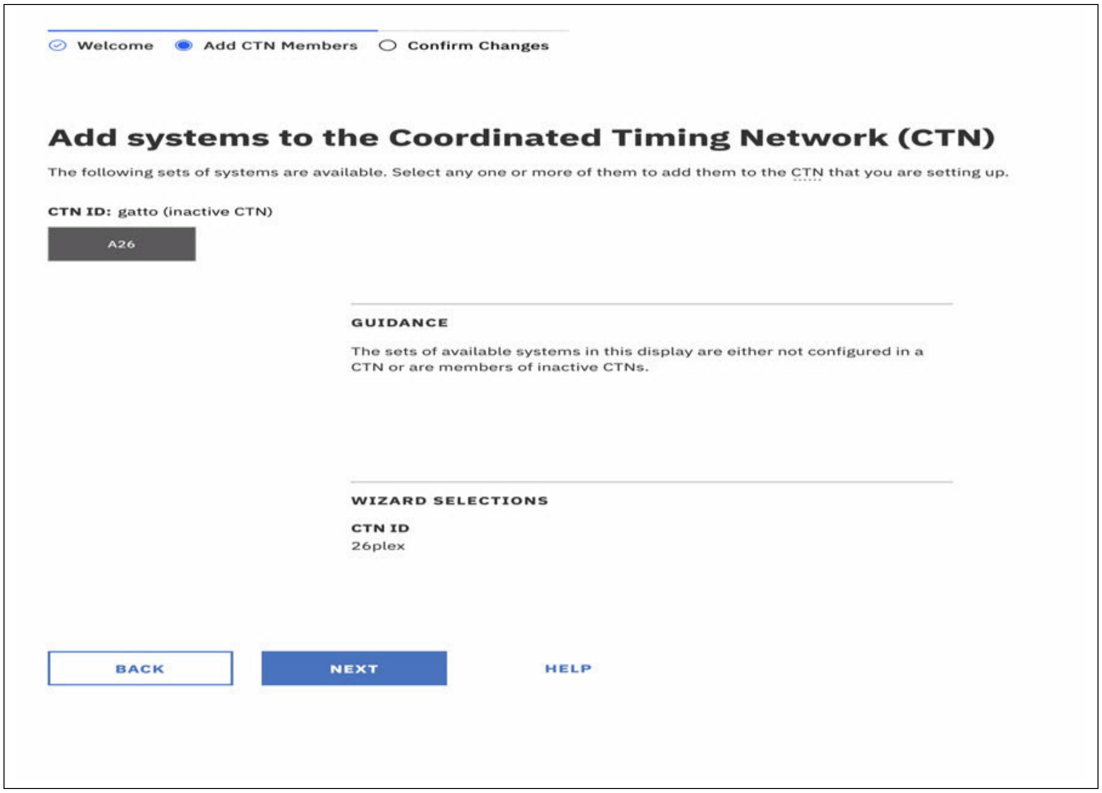

Figure 8-8 Add system to CTN

3. The Confirm Changes window opens and shows you, how and on which STRATUM level the new CPC will be added as shown in Figure 4. Click **APPLY** to continue.

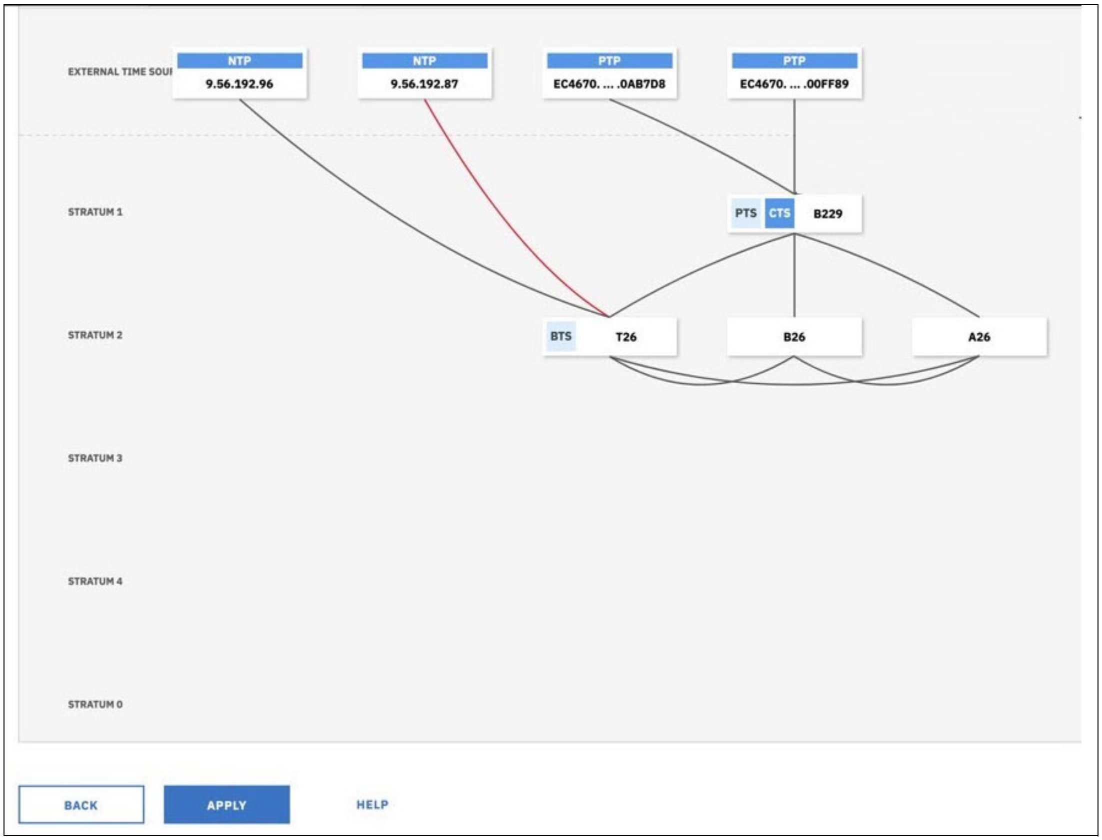

Figure 8-9 APPLY adding a CPC to a CTN

4. You see the Local CTN ID change confirmation for the selected CPC. After you select **APPLY**, the Complete message window opens. Click **Close** to complete the operation.

8.4.3 Verifying the new CTN configuration

When the new CTN configuration completes, the Manage System Time initial window opens again, showing the changed CTN.

Check that the following items are defined correctly in the window:

- Status
- CTN ID
- Time / Date / Time zone
- Stratum Level
- STP Role

You can also check the status of STP by running the z/OS **D ETR** command (Example 8-1).

Example 8-1 D ETR display command

```
D ETR
IEA386I 11.48.43 TIMING STATUS 977
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = PCC-TIME     ETS ID = DNTP (Direct attachment using NTP)
  ETS Timing Precision = 2
```

```
THE STRATUM 1 NODE ID = 003932.AGZ.IBM.02.000000087F28
THIS IS THE ARBITER SERVER
NUMBER OF USABLE TIMING LINKS =  7
```

For more information, see *IBM Z Server Time Protocol Guide*, SG24-8480.

9

Defining Coupling Facility links

This chapter describes the coupling connectivity options and the Parallel Sysplex clustering enhancements that are available on IBM z17 system. Coupling link configuration examples for Parallel Sysplex and Server Time Protocol (STP) are also shown.

For more information about the enhancements that were made in Coupling Facility (CF) level 26, see *IBM z17 (9175) Technical Guide*, SG24-8579.

Note: The examples that are shown in this chapter are based on the IBM z17 ME1 (9175).

This chapter includes the following topics:

- Coupling connectivity options for Parallel Sysplex on IBM z17
- Defining Coupling Express3 Long Reach
- Defining Integrated Coupling Adapter Short Reach
- Defining an STP timing-only link by using ICA SR
- CF LPAR setup and Coupling Facility Control Code Level 26
- Dynamic I/O for Stand-alone Coupling Facility

9.1 Coupling connectivity options for Parallel Sysplex on IBM z17

Coupling connectivity for Parallel Sysplex on IBM z17 can use Coupling Express3 Long Reach (CE3 LR) and Integrated Coupling Adapter Short Reach 2.0 (ICA SR2.0). The ICA SR feature is designed to support distances of up to 150 m. The CE3 LR feature supports distances up to 10 km unrepeated between systems, and up to 100 km with qualified Dense Wavelength Division Multiplexer (DWDM). The available options for coupling links on IBM z17 are shown in Figure 9-1.

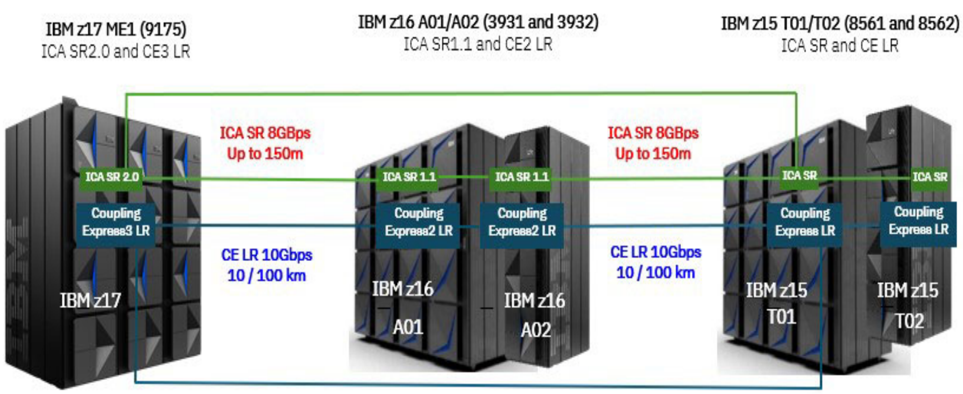

Figure 9-1 IBM z17 coupling connectivity

Internal Coupling (IC) links are supported and used for internal communication between logical partitions (LPARs) on the same central processor complex (CPC) running Coupling Facilities (CFs) and z/OS images. The connection is emulated in the Licensed Internal Code (LIC) and provides fast and secure memory-to-memory communications between LPARs within a single system. No physical cabling is required. For IC, an ICP type channel path ID (CHPID) is used.

Note: ICA SR, ICA SR 1.1, CE LR and CE2 LR are not available on IBM z17, neither as new build or carry-forward. Only ICA SR2.0 and CE3 LR external coupling links are supported.

9.1.1 Preparing to define Coupling Facility links

A good starting point for implementing coupling links is accurate and current documentation that clearly illustrates all connections that are needed for the new or upgraded CPC.

When installing coupling links, ensure that you ordered enough ports to support your configuration with physical feature redundancy. Your Parallel Sysplex should be configured for the highest possible availability.

Parallel Sysplex failure independence is a function of a z/OS to CF relationship and the removal of single points of failure. For example, all connections to a structure on a Stand-alone Coupling Facility (SACF) are failure-independent. With an Internal Coupling Facility (ICF), all connections from z/OS images on the same footprint are failure-dependent.

For more information, see *Coupling Facility Configuration Options*, ZSW01971, which can be found at the IBM downloads website.

Evaluate whether the configuration includes any channel features that are not supported on IBM z17. The configuration should be reviewed for any channel types that cannot be carried forward or connected to the IBM z17.

Another important point is to ensure that all CPCs are connected to the IBM z17 by using coupling links as follows: Only N, N-1, and N-2 IBM Z generations can coexist in the same Parallel Sysplex or STP Coordinated Timing Network (CTN). For example, IBM z17 provides coupling connectivity back to IBM z16 A01/A02, IBM z15 T01/T02 systems only through ICA SR2.0 and CE3-LR features.

Note: Deactivate any coupling links on other connected systems before an upgrade, or you might experience configuration errors.

If an IBM z17 plays a CTN role (Preferred Time Server (PTS), Backup Time Server (BTS), or Arbiter), then the other CTN role-playing CPCs must have coupling connectivity to the IBM z17 (N, N-1, and N-2).

If coupling links are connected across sites by using Dense Wavelength Division Multiplexing (DWDM), verify whether the DWDM equipment that you plan to use supports the respective coupling link technology and is qualified for the STP.

Note: Do not use DWDM equipment that is not qualified by IBM to transport STP information.

For more information about the supported coupling link features, see:
IBM Z Connectivity Handbook, SG24-5444.

9.2 Defining Coupling Express3 Long Reach

This section describes the implementation of the CE3-LR feature. The definition of this CF link is part of the activity that is called *Define CF/STP link*, which is shown in Figure 1-3 on page 5.

On IBM z17, the long-reach coupling link adapter hardware is upgraded from Coupling Express2 LR to a new Coupling Express3 LR CX6-LX adapter. Maximum 32 CE3-LR adapters per CPC.

This new CE3 LR adapter will support 2 varieties of optics:

- 10Gb (as today);
- A new 25Gb option for higher bandwidth and higher capacity/potential throughput.

When configured to use 10Gb optics, the adapter will remain compatible with existing CE LR and CE2 LR (CL5) links on previous machines like z16 and z15, and will be represented as a CL5 coupling link type.

> **Note:** CL5 only connects to CL5; Use for connections back to previous machines using CL5, or for connections to other IBM z17 machines still using CL5.

Coupling Express3 LR (CE LR) coexistence/migration is critical:

- IBM z17 CE3-LR 10Gb links connect to existing CE LR links on IBM z15 and CE2 LR on z16 by using the same LR cabling.
- IBM z17 introduces new support for Coupling Express3 LR 25G Coupling Link adapters to provide long-distance coupling link connectivity with 25 gigabit bandwidth (CL6), in addition to supporting the existing 10 gigabit bandwidth using Coupling Express3 LR 10G (CL5).
- Customers should plan their migration off of CL5 and onto CL6 Coupling Links accordingly.

> **Important:** The machine after z17 is planned to be the last IBM Z hardware generation which will support the 10 gigabit bandwidth (CL5) coupling links. The subsequent hardware generation after that is planned to support only 25 gigabit bandwidth CL6 coupling links for long-distance coupling link connectivity.

When configured to use 25Gb optics, the adapter will be incompatible with existing CE LR and CE2 LR (CL5) links on previous machines and can only connect to another 25Gb CE3 LR adapter. Because of this incompatibility, we will define and use a new CL6 coupling link type for this.

- CL6 only connects to CL6; Use ONLY for connections to other IBM z17 machines using CL6;
- Other than the higher bandwidth and new link type, CL6 looks and behaves much the same as CL5, except that in the HCD definitions:
  - CL5 defaults to definition of 8 subchannels/devices per CHPID, with an option for 32;
  - CL6 is the other way around: CL6 defaults to definition of 32 subchannels/devices per CHPID, with an option for 8.
- Also, because of its higher bandwidth, CL6 operates in a higher path selection “preference tier” than CL5; other things being equal, selection of CL6 CHPIDs is preferred over CL5.

9.2.1 CE3 LR 10Gb (CL5) connections:

In this example, we show the connection of a coupling link between an IBM z17 and an IBM z16 or IBM z15 by using CHPID type CL5 (see Figure 9-2).

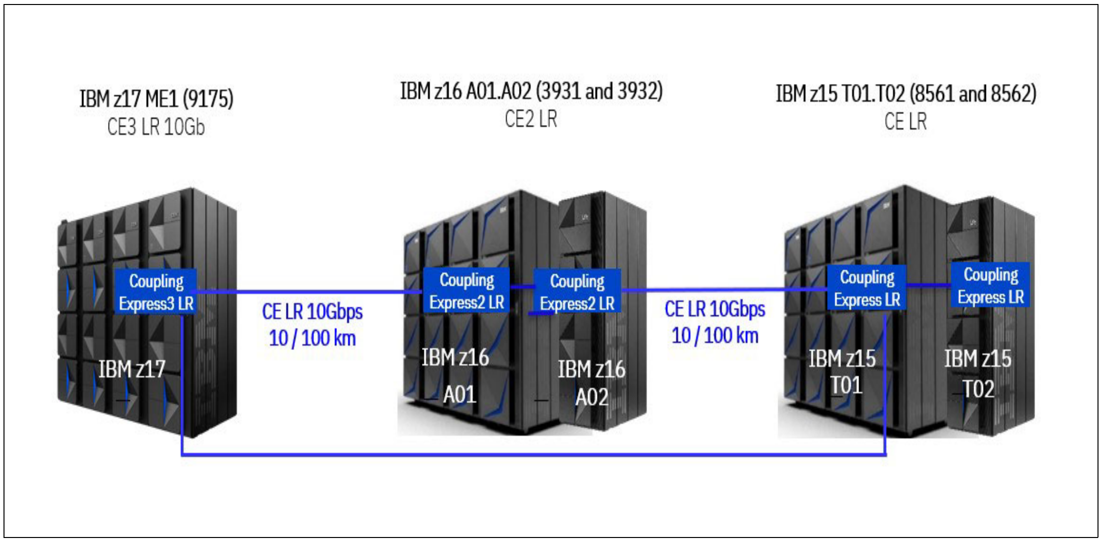

Figure 9-2 CF link connection that uses CL5

CE3-LR 10Gb CL5 is defined in an input/output configuration data set (IOCDS). Even though this feature is a Peripheral Component Interconnect Express (PCIe) feature, a physical channel ID (PCHID) is used instead of an adapter ID (AID) to identify the physical card.

Example 9-1 shows a sample extract of the input/output configuration program (IOCP) to define the new CHPID Type CL5 on the IBM z16 that connects to an IBM z17.

Example 9-1 IOCP definitions for CHPID Type CL5 on an IBM z16

```
ID     ..                                                       *
       ..                                                       *
       SYSTEM=(3931,1),LSYSTEM=PAVO,                            *
       TOK=('PAVO',008003009FB89175143413700125107F00000000,000*
       00000,'25-04-04','08:09:52','SYS9','IODF79')
RESOURCE PARTITION=((CSS(0),(PAVO0A,A),..                       *
       PAVO0C,C),(PAVO0D,D),(PAVO0E,E),(PAVO0F,F),(PAVO01,1),(P*
       ..                                                       *
       (CSS(3),(PAVO3A,A),(PAVO3B,B),(PAVO3C,C),(PAVO3D,D),(PAV*
       03E,E),..                                                *
CHPID PATH=(CSS(3),88),SHARED,PARTITION=((PAVO3E),(=)),         *
       CPATH=(CSS(3),88),CSYSTEM=MENSA,PORT=1,PCHID=40C,        *
       TYPE=CL5
```

Example 9-2 shows a sample extract of the corresponding IOCP definition for a connecting CHPID Type CL5 on an IBM z17.

Example 9-2 IOCP definitions for CHPID Type CL5 on an IBM z17 system

```
ID     ..                                                       *
       ..
       SYSTEM=(9175,1),LSYSTEM=MENSA,                           *
       TOK=('MENSA',008003009FB89175143413700125107F00000000,00*
```

```
      000000,'25-04-17','14:34:13','SYS9','IODF79')
RESOURCE PARTITION=((CSS(0),(MENSA0A,A),(MENSA0B,B),(MENSA0C,C*
      ),(MENSA0D,D),(MENSA01,1),(MENSA02,2),(MENSA03,3),(MENSA*
      ..
      F,F),(MENSA31,1),(MENSA32,2),(MENSA33,3),(MENSA34,4),(ME*
      NSA35,5),...                                            *
CHPID PATH=(CSS(3,5),88),SHARED,                              *
      PARTITION=((CSS(3),(MENSA32,MENSA3E),(=)),(CSS(5),(MENSA*
      51),(=))),CPATH=(CSS(3),88),CSYSTEM=PAVO,PORT=1,        *
      PCHID=17C,TYPE=CL5
```

Note: When you connect CF sender and receiver channel paths, or CF peer channel paths, the Hardware Configuration Definition (HCD) proposes a CF control unit (CU), and device numbers that must be defined for a CF sender channel (CF receiver channels do not require CUs and devices to be defined).

For more information about how to define CHPID Type CL5 in HCD, see 14.2.4, “Defining CL5 and CL6 CHPIDs” on page 312.

9.2.2 CE3 LR 10Gb (CL5) verifying the configuration:

After you activate the new configuration with the new CE3-LR (10Gb) CF links and all the cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the Support Element (SE) windows:

- Checking the status by using z/OS commands.

  For example, if you are interested in checking the status of CHPID 88, run the **D M=CHP(88)** command, as shown in Example 9-3.

Example 9-3 Displaying the status of CHPID 88

```
D M=CHP(88)
IEE174I 09.27.10 DISPLAY M 378
CHPID 88:  TYPE=34, DESC=COUPLING OVER ROCE 10G, ONLINE
COUPLING FACILITY 003931.IBM.02.000000071A08
                  PARTITION: 3E  CPCID: 00
NAMED CF78        CONTROL UNIT ID: FFF0

PATH        PHYSICAL              LOGICAL  CHANNEL TYPE     CAID POR
88 / 040C   ONLINE                ONLINE   CL5 10GbE-RoCE   017C  01

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   96   IN USE:   72   NOT USING:    0   NOT USABLE:   24
 OPERATIONAL DEVICES / SUBCHANNELS:
     FD2E / 170D     FD2F / 170E     FD30 / 170F     FD31 / 1710
     FD32 / 1711     FD33 / 1712     FD34 / 1713     FD35 / 1714
     FD36 / 1715     FD37 / 1716     FD38 / 1717     FD39 / 1718
     FD3A / 1719     FD3B / 171A     FD3C / 171B     FD3D / 171C
     FD3E / 171D     FD3F / 171E     FD40 / 171F     FD41 / 1720
     FD42 / 1721     FD43 / 1722     FD44 / 1723     FD45 / 1724
     FD46 / 1725     FD47 / 1726     FD48 / 1727     FD49 / 1728
     FD4A / 1729     FD4B / 172A     FD4C / 172B     FD4D / 172C
     FDF5 / 172D     FDF6 / 172E     FDF7 / 172F     FDF8 / 1730
     FDF9 / 1731     FDFA / 1732     FDFB / 1733     FDFC / 1734
     FDFD / 1735     FDFE / 1736     FDFF / 1737     FE00 / 1738
```

```
FE01 / 1739    FE02 / 173A    FE03 / 173B    FE04 / 173C
FE05 / 173D    FE06 / 173E    FE07 / 173F    FE08 / 1740
FE09 / 1741    FE0A / 1742    FE0B / 1743    FE0C / 1744
FE0D / 1745    FE0E / 1746    FE0F / 1747    FE10 / 1748
FE11 / 1749    FE12 / 174A    FE13 / 174B    FE14 / 174C
FE35 / 174D    FE36 / 174E    FE37 / 174F    FE38 / 1750
FE39 / 1751    FE3A / 1752    FE3B / 1753    FE3C / 1754
FE3D / 1755    FE3E / 1756    FE3F / 1757    FE40 / 1758
FE41 / 1759    FE42 / 175A    FE43 / 175B    FE44 / 175C
FE45 / 175D    FE46 / 175E    FE47 / 175F    FE48 / 1760
FE49 / 1761    FE4A / 1762    FE4B / 1763    FE4C / 1764
FE4D / 1765    FE4E / 1766    FE4F / 1767    FE50 / 1768
FE51 / 1769    FE52 / 176A    FE53 / 176B    FE54 / 176C
```

- Checking the status by using the SE windows by completing the following steps:
  a. From the Hardware Management Console (HMC), select the CPC (under Systems Management) where the CHPID/PCHID that you want to verify is, and click **Single Object Operations** (under Recovery task options).
  b. On the SE, select the same CPC and click **Channels**, as shown in Figure 9-3.

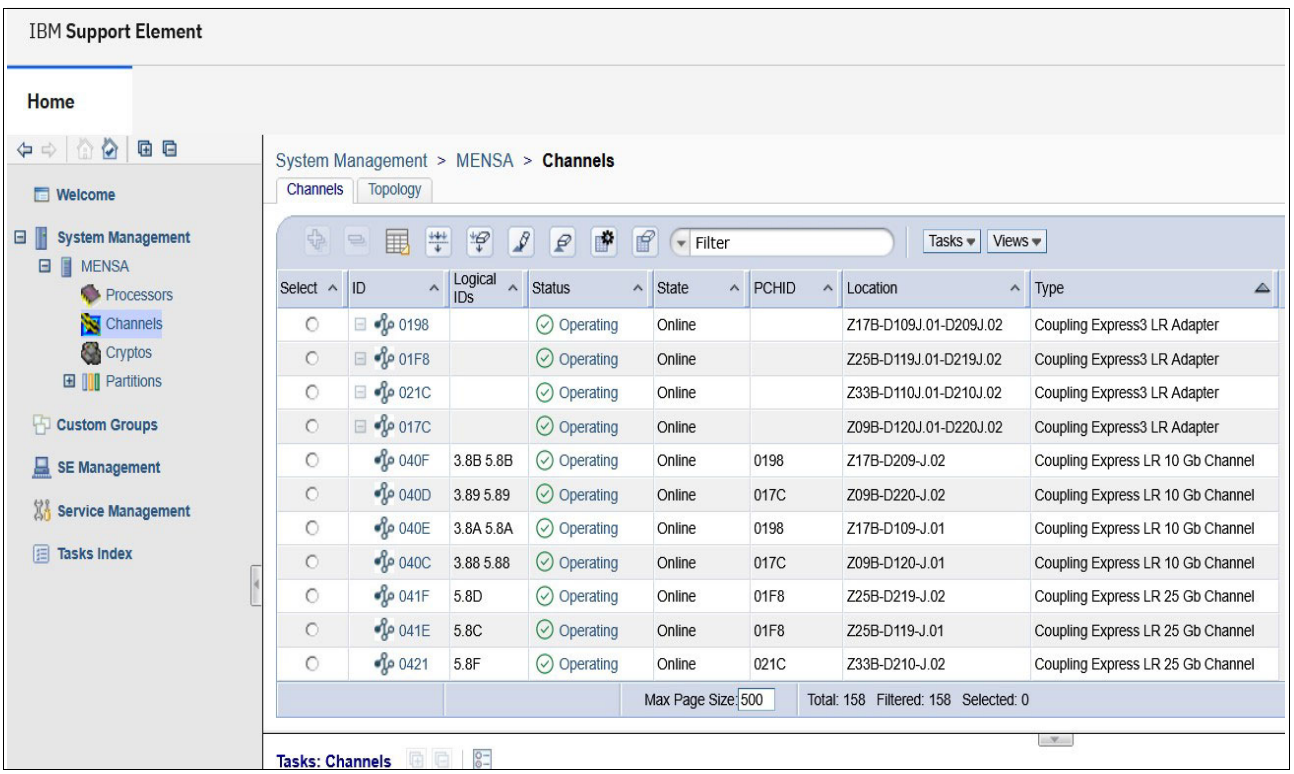

Figure 9-3 SE Systems Management and channels

  c. Look for the PCHID that you are interested in checking the status of. The result looks like what is shown in Figure 9-4 on page 172.

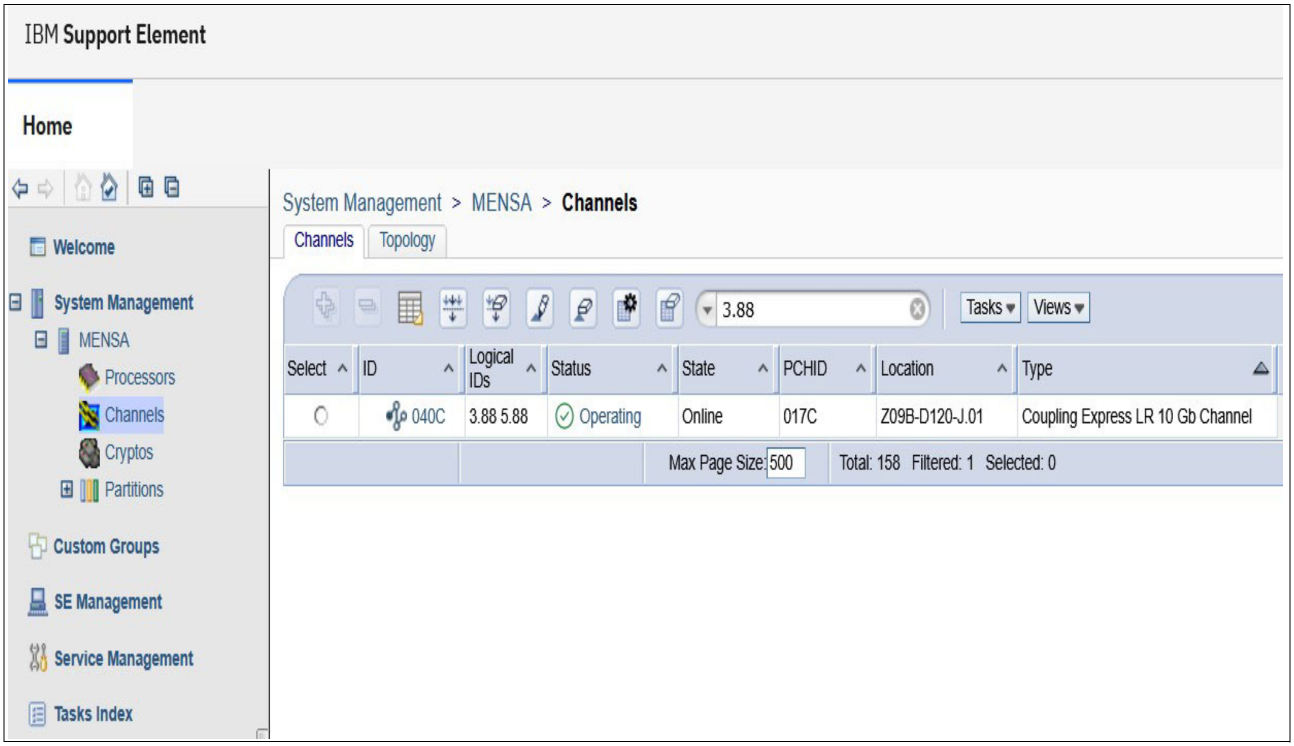

Figure 9-4 Verifying channel CL5 by using the CPC view

For more information about the PCHID, click the PCHID to show its details, as shown in Figure 9-5 on page 173.

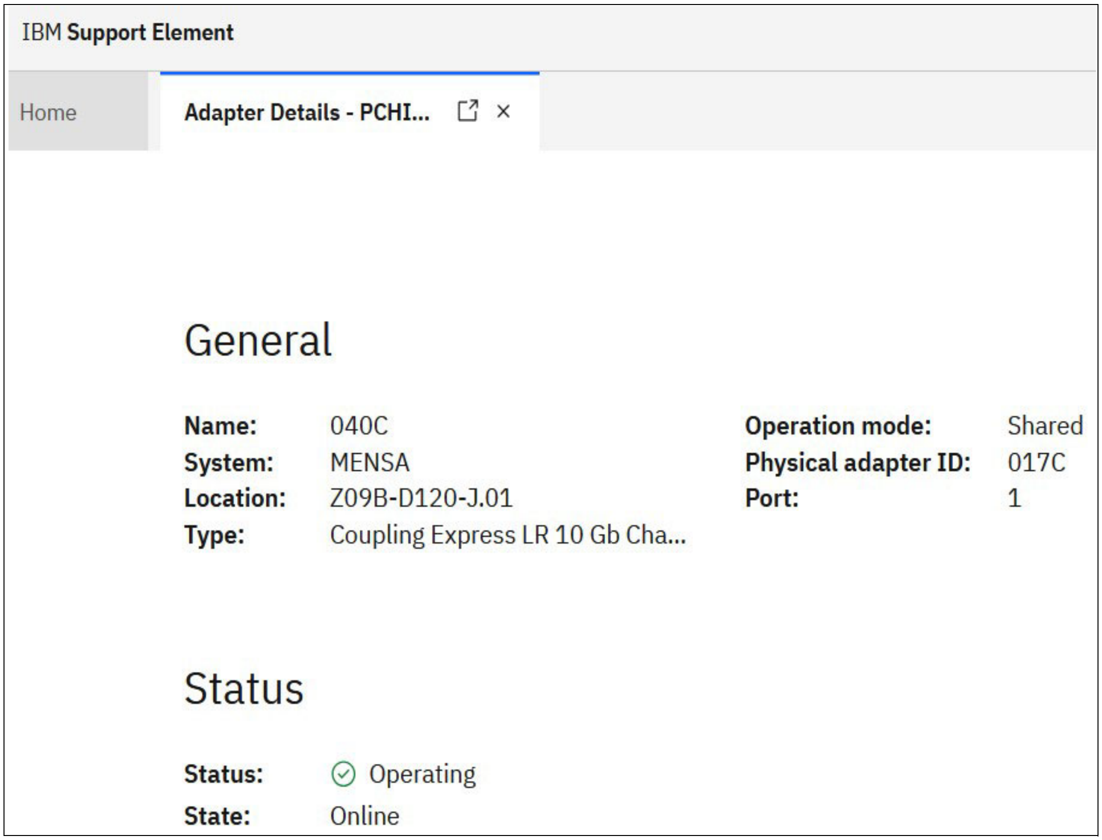

Figure 9-5 CL5 PCHID details

You can also check the status from the LPAR view. Select the LPAR that you want to check the CHPID status of, and select the channels option under that LPAR. Now, you can look for the CHPID and check the status, as shown in Figure 9-6 on page 174.

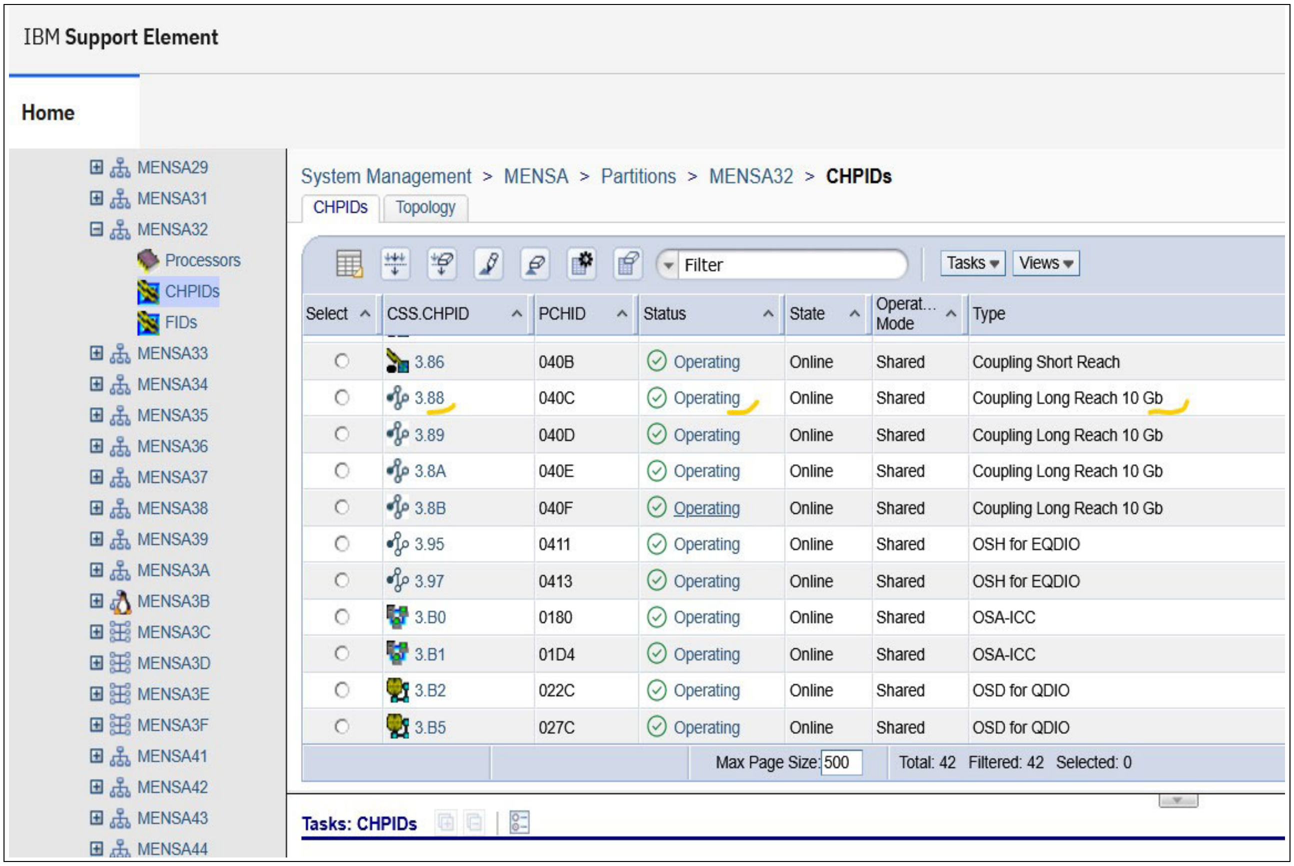

Figure 9-6 SE Verify channel LPAR view

9.2.3 CE3 LR 25Gb (CL6) connections:

In this example, we show the connection of a coupling link between two IBM z17 by using CHPID type CL6 (see Figure 9-7).

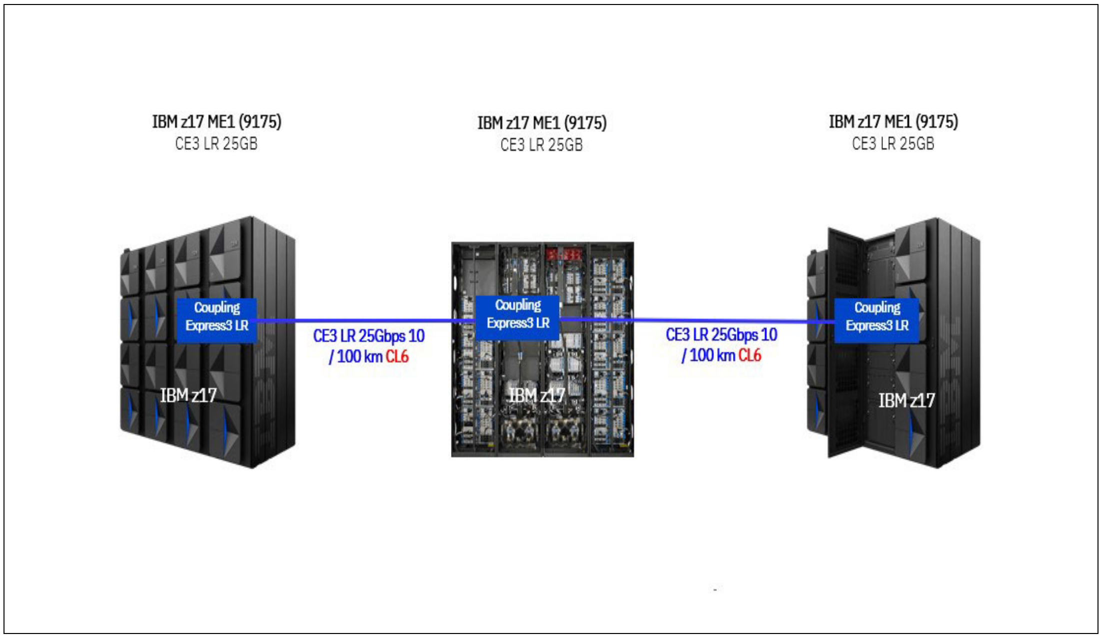

Figure 9-7 CF link connection that uses CL6

External connectivity between z17s allows the new 25Gb Coupling Express 3 Long Reach links. Those connections are defined as a new chpid type called CL6 instead of the CL5 used for the CE3 LR 10 Gb. The CL6 channels are defined like the CL5 channels but with 32 subchannels/devices per CHPID as default, with an option for 8.

Example 9-4 shows a sample extract of the corresponding IOCP definition for a connecting CHPID Type CL5 on an IBM z17.

Example 9-4 IOCP definitions for CHPID Type CL6 on a z17 system

```
ID     MSG1='IODF79',MSG2='SYS9.IODF79 - 2025-04-17 14:34',     *
       SYSTEM=(9175,1),LSYSTEM=MENSA,                           *
       TOK=('MENSA',008003009FB89175143413700125107F00000000,00*
       000000,'25-04-17','14:34:13','SYS9','IODF79')
       .....
RESOURCE PARTITION=((CSS(3),(MENSA3A,A),(MENSA3B,B),(MENSA3C,C*
       ),(MENSA3D,D),(MENSA3E,E),(MENSA3F,F),(MENSA31,1),(MENSA*
       32,2),(MENSA33,3),(MENSA34,4),(MENSA35,5),(MENSA36,6),(M*
       ENSA37,7),(MENSA38,8),(MENSA39,9))),                     *
       .....
CHPID PATH=(CSS(3),C0),SHARED,PARTITION=((MENSA32),(=)),        *
       TPATH=((CSS(3),MENSA,C4),(CSS(3),MENSA,C0,FFE8,FD0E,32))*
       ,DESC='LR 25G',PORT=1,PCHID=1F8,TYPE=CL6
CHPID PATH=(CSS(3),C1),SHARED,PARTITION=((MENSA32),(=)),        *
       TPATH=((CSS(3),MENSA,C5),(CSS(3),MENSA,C1,FFE8,FCEE,32))*
       ,DESC='LR 25G',PORT=2,PCHID=1F8,TYPE=CL6
CHPID PATH=(CSS(3),C2),SHARED,PARTITION=((MENSA32),(=)),        *
       TPATH=((CSS(3),MENSA,C6),(CSS(3),MENSA,C2,FFE9,FC7E,32))*
       ,DESC='LR 25G',PORT=1,PCHID=1F8,TYPE=CL6
CHPID PATH=(CSS(3),C3),SHARED,PARTITION=((MENSA32),(=)),        *
       TPATH=((CSS(3),MENSA,C7),(CSS(3),MENSA,C3,FFE9,FC5E,32))*
       ,DESC='LR 25G',PORT=2,PCHID=1F8,TYPE=CL6
CHPID PATH=(CSS(3),C4),SHARED,PARTITION=((MENSA3E),(=)),        *
```

```
      TPATH=((CSS(3),MENSA,C0,FFE8,FD0E,32),(CSS(3),MENSA,C4))*
      ,DESC='LR 25G',PORT=1,PCHID=21C,TYPE=CL6
CHPID PATH=(CSS(3),C5),SHARED,PARTITION=((MENSA3E),(=)),      *
      TPATH=((CSS(3),MENSA,C1,FFE8,FCEE,32),(CSS(3),MENSA,C5))*
      ,DESC='LR 25G',PORT=2,PCHID=21C,TYPE=CL6
CHPID PATH=(CSS(3),C6),SHARED,PARTITION=((MENSA3F),(=)),      *
      TPATH=((CSS(3),MENSA,C2,FFE9,FC7E,32),(CSS(3),MENSA,C6))*
      ,DESC='LR 25G',PORT=1,PCHID=21C,TYPE=CL6
CHPID PATH=(CSS(3),C7),SHARED,PARTITION=((MENSA3F),(=)),      *
      TPATH=((CSS(3),MENSA,C3,FFE9,FC5E,32),(CSS(3),MENSA,C7))*
      ,DESC='LR 25G',PORT=2,PCHID=21C,TYPE=CL6
```

Considering the same configuration above, on the HCD under CF Channel Path Connectivity List panel you can see the cflinks connected and there are some CL6 as shown on Example 9-5.

Example 9-5 CF Channel Path Connectivity List

```
---------------------------------------------------------------------------
                     CF Channel Path Connectivity List            Row 1 of 10
 Command ===> ____________________________________________ Scroll ===> PAGE

 Select one or more channel paths, then press Enter.

 Source processor ID . . . . . : MENSA     Mensa
 Source channel subsystem ID . : 5
 Source partition name . . . . : *

   ----------Source----------   -----------Destination-----------   -CU-  -#-
 / CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
 _ 85  04/2  Y  CS5  SHR  N     VELA.2     D3  04/2  Y  CS5  SHR    CFP   32
 _ 87  10/2  Y  CS5  SHR  N     VELA.2     D7  06/2  Y  CS5  SHR    CFP   32
 _ 88  17C/1 Y  CL5  SPAN N     PAVO.3     88  1B8/1 Y  CL5  SHR    CFP   32
 _ 89  17C/2 Y  CL5  SPAN N     PAVO.3     8B  1DC/2 Y  CL5  SHR    CFP   32
 _ 8A  198/1 Y  CL5  SPAN N     PAVO.3     8A  1DC/1 Y  CL5  SHR    CFP   32
 _ 8B  198/2 Y  CL5  SPAN N     PAVO.3     89  1B8/2 Y  CL5  SHR    CFP   32
 _ 8C  1F8/1 N  CL6  SHR  N     MENSA.5    8E  21C/1 Y  CL6  SHR    CFP   32
 _ 8D  1F8/2 N  CL6  SHR  N     MENSA.5    8F  21C/2 Y  CL6  SHR    CFP   32
 _ 8E  21C/1 Y  CL6  SHR  N     MENSA.5    8C  1F8/1 N  CL6  SHR    CFP   32
 _ 8F  21C/2 Y  CL6  SHR  N     MENSA.5    8D  1F8/2 N  CL6  SHR    CFP   32
 ****************************** Bottom of data ******************************
```

9.2.4 CE3 LR 25Gb (CL6) verifying the configuration:

After you activate the new configuration with the new CE3-LR (25Gb) CF links and all the cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the Support Element (SE) windows:

- Checking the status by using z/OS commands.

For example, if you are interested in checking the status of CHPID 88, run the **`D M=CHP(C0)`** command, as shown in Example 9-6.

Example 9-6 Displaying the status of CHPID C0

```
D M=CHP(C0)
IEE174I 15.21.52 DISPLAY M 549
CHPID C0:  TYPE=37, DESC=COUPLING OVER ROCE 25G, ONLINE
COUPLING FACILITY 009175.IBM.02.0000000B9FB8
                  PARTITION: 3E  CPCID: 00
NAMED CF76        CONTROL UNIT ID: FFE8

PATH        PHYSICAL              LOGICAL  CHANNEL TYPE     CAID PORT
C0 / 041E   ONLINE                ONLINE   CL6 25GbE-RoCE   01F8  01

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   80   IN USE:   80   NOT USING:    0   NOT USABLE:    0
 OPERATIONAL DEVICES / SUBCHANNELS:
     FCEE / 1839     FCEF / 183A     FCF0 / 183B     FCF1 / 183C
     FCF2 / 183D     FCF3 / 183E     FCF4 / 183F     FCF5 / 1840
     FCF6 / 1841     FCF7 / 1842     FCF8 / 1843     FCF9 / 1844
     FCFA / 1845     FCFB / 1846     FCFC / 1847     FCFD / 1848
     FCFE / 1849     FCFF / 184A     FD00 / 184B     FD01 / 184C
     FD02 / 184D     FD03 / 184E     FD04 / 184F     FD05 / 1850
     FD06 / 1851     FD07 / 1852     FD08 / 1853     FD09 / 1854
     FD0A / 1855     FD0B / 1856     FD0C / 1857     FD0D / 1858
     FD0E / 1859     FD0F / 185A     FD10 / 185B     FD11 / 185C
     FD12 / 185D     FD13 / 185E     FD14 / 185F     FD15 / 1860
     FD16 / 1861     FD17 / 1862     FD18 / 1863     FD19 / 1864
     FD1A / 1865     FD1B / 1866     FD1C / 1867     FD1D / 1868
     FD1E / 1869     FD1F / 186A     FD20 / 186B     FD21 / 186C
     FD22 / 186D     FD23 / 186E     FD24 / 186F     FD25 / 1870
     FD26 / 1871     FD27 / 1872     FD28 / 1873     FD29 / 1874
     FD2A / 1875     FD2B / 1876     FD2C / 1877     FD2D / 1878
     FF80 / 17CD     FF81 / 17CE     FF82 / 17CF     FF83 / 17D0
     FF84 / 17D1     FF85 / 17D2     FF86 / 17D3     FF87 / 17D4
     FFF0 / 17D5     FFF1 / 17D6     FFF2 / 17D7     FFF3 / 17D8
     FFF4 / 17D9     FFF5 / 17DA     FFF6 / 17DB     FFF7 / 17DC
```

- Checking the status by using the SE windows by completing the following steps:
  a. From the Hardware Management Console (HMC), select the CPC (under Systems Management) where the CHPID/PCHID that you want to verify is, and click **Single Object Operations** (under Recovery task options).
  b. On the SE, select the same CPC and click **Channels**, as shown in Figure 9-8 on page 178.

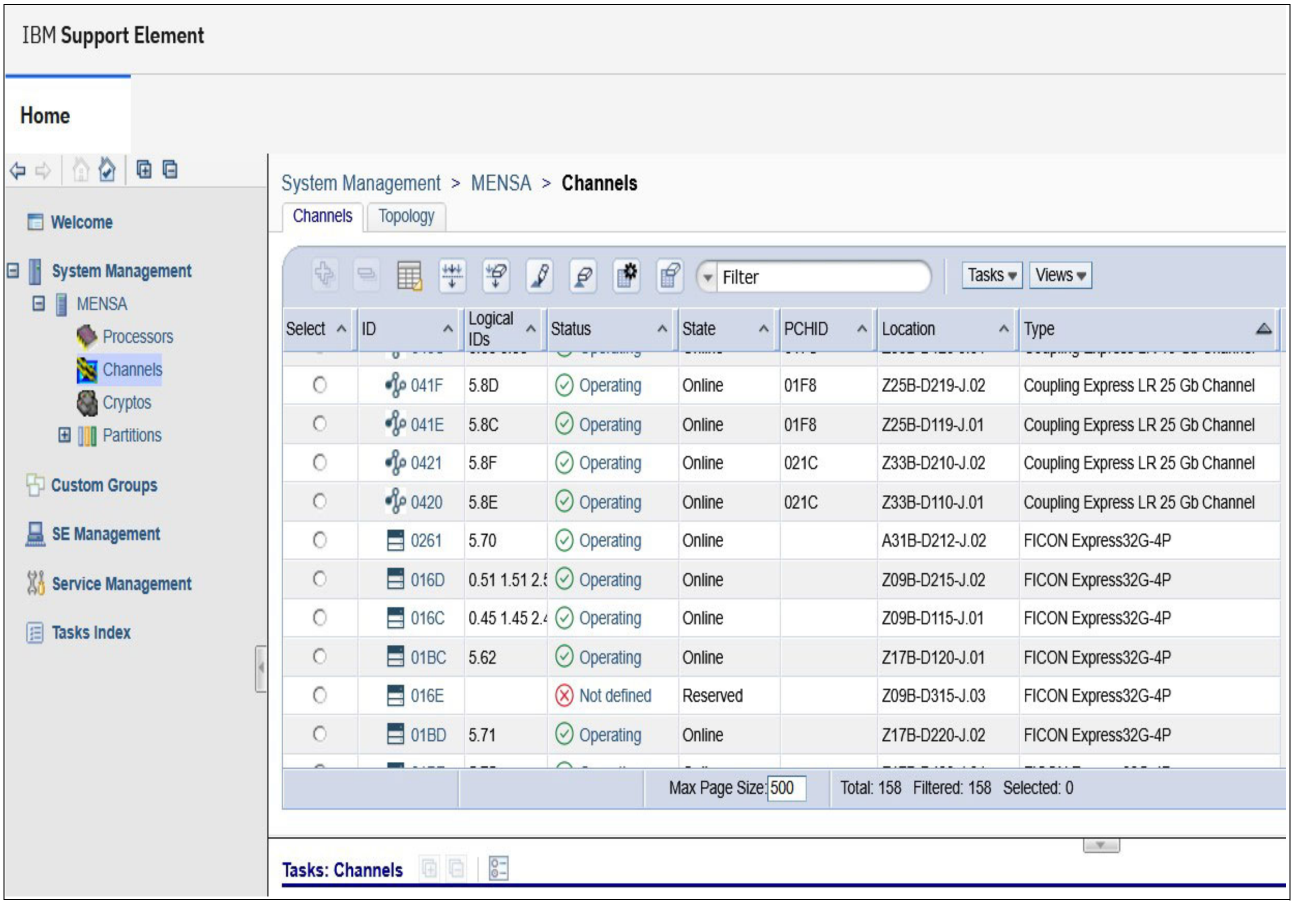
Figure 9-8 CL6 on SE Systems Management and channels

9.3 Defining Integrated Coupling Adapter Short Reach

This section describes the implementation of coupling links by using the ICA-SR 2.0 feature. The definition of this CF link is part of the activity that is called *Define CF/STP link*, which is shown in Figure 1-3 on page 5.

The ICA-SR adapter on IBM z17 is changing from ICA SR 1.1 to a new ICA SR 2.0 adapter.

> **Note:** ICA SR and ICA SR 1.1 will not be available on IBM z17, neither as new build or carry-forward.

- ICA-SR on IBM z17 summary:
  - Follow-on to the ICA SR 1.0 and ICA SR 1.1
    - Fully Compatible with ICA SR and ICA SR 1.1;
  - Improved Large block write command completion without DREQ-less;
  - Per-message protocol selection (vs. full port);
  - Fan-out in the CPC drawer, 2-ports per fan-out, 150m;
  - Remains as link type CS5;
  - 8 primary send buffers / CHPID;

– 4 logical CHPIDs / port;
– Maximum 48 adapters (96 ports) per CEC.

In this example, we show how to define a coupling link between an IBM z17 system and an IBM z15 or z16 system by using CHPID type CS5 (see Figure 9-9).

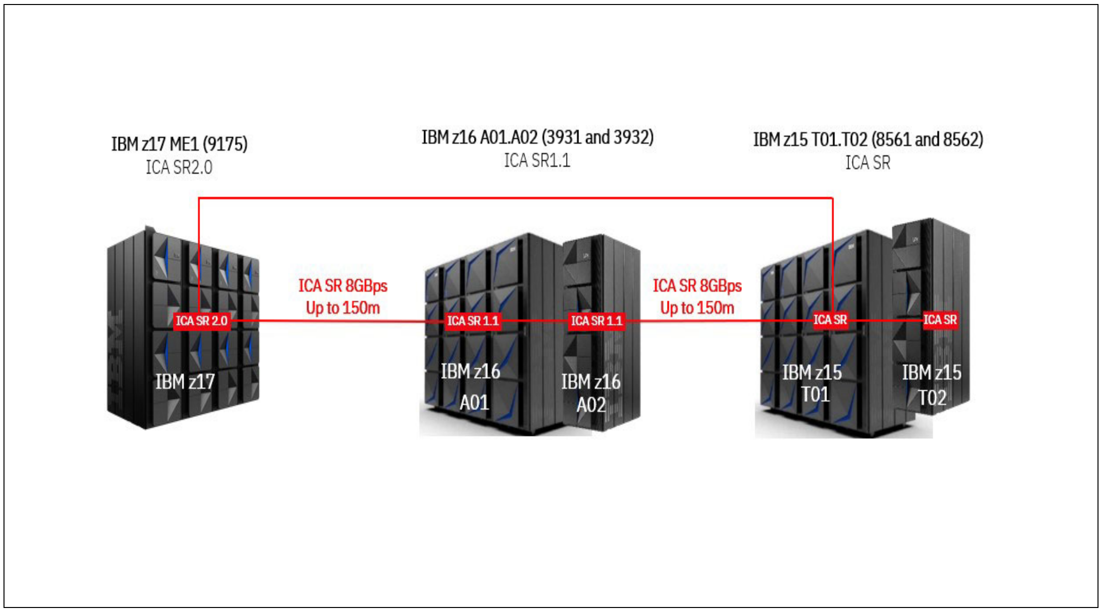

Figure 9-9 CF link CS5 connection from an IBM z17 to an IBM z15 or z16

The ICA SR is defined in an IOCDS-like PSIFB by using an AID to identify the physical card. Example 9-7 shows a sample extract of the corresponding IOCP definition for a connecting CHPID Type CS5 on an IBM z17.

Example 9-7 IOCP definitions for CHPID Type CS5 on an IBM z17

```
ID     MSG1='IODF79',MSG2='SYS9.IODF79 - 2025-04-17 14:34',      *
       SYSTEM=(9175,1),LSYSTEM=MENSA,                            *
       TOK=('MENSA',008003009FB89175143413700125107F00000000,00*
       000000,'25-04-17','14:34:13','SYS9','IODF79')
RESOURCE PARTITION=((CSS(0),(MENSA0A,A),(MENSA0B,B),(MENSA0C,C*
       ),(MENSA0D,D),(MENSA01,1),(MENSA02,2),(MENSA03,3),(MENSA*
       ....
       F,F),(MENSA31,1),(MENSA32,2),(MENSA33,3),(MENSA34,4),(ME*
       NSA35,5),(MENSA36,6),(MENSA37,7),(MENSA38,8),(MENSA39,9)*
       ....
CHPID PATH=(CSS(3),84),SHARED,                                   *
       PARTITION=((MENSA3E,MENSA32),(=)),CPATH=(CSS(3),85),      *
       CSYSTEM=PAVO,AID=04,PORT=1,TYPE=CS5
CHPID PATH=(CSS(3),86),SHARED,                                   *
       PARTITION=((MENSA3F,MENSA32),(=)),CPATH=(CSS(3),81),      *
       CSYSTEM=PAVO,AID=10,PORT=1,TYPE=CS5
```

For more information about how to define CHPID Type CS5 in the HCD, see 14.2.3, “Defining a Coupling Facility link with CS5 CHPIDs” on page 309.

9.3.1 ICA SR: Verifying the configuration

After activating the new configuration with the ICA SR CF links and all cables are connected, verify whether the CHPIDs are online and operating by using a z/OS command or the SE windows:

- Checking the status by using the z/OS command.

  For example, if you are interested in checking the status of CHPID 84, run a **D M=CHP(84)** command, as shown in Example 9-8.

Example 9-8 Displaying the status of CHPID 84

```
D M=CHP(84)
IEE174I 10.44.29 DISPLAY M 287
CHPID 84:  TYPE=33, DESC=COUPLING OVER PCIE, ONLINE
COUPLING FACILITY 003931.IBM.02.000000071A08
                  PARTITION: 3F  CPCID: 00
NAMED CF79        CONTROL UNIT ID: FFFB

PATH        PHYSICAL              LOGICAL  CHANNEL TYPE      CAID PORT
84 / 040A   ONLINE                ONLINE   CS5 8X-PCIE3      0004  01

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   96   IN USE:   72   NOT USING:    0   NOT USABLE:   24
 OPERATIONAL DEVICES / SUBCHANNELS:
     FD75 / 176D     FD76 / 176E     FD77 / 176F     FD78 / 1770
     FD79 / 1771     FD7A / 1772     FD7B / 1773     FD7C / 1774
     FD7D / 1775     FD7E / 1776     FD7F / 1777     FD80 / 1778
     FD81 / 1779     FD82 / 177A     FD83 / 177B     FD84 / 177C
     FD85 / 177D     FD86 / 177E     FD87 / 177F     FD88 / 1780
     FD89 / 1781     FD8A / 1782     FD8B / 1783     FD8C / 1784
     FD8D / 1785     FD8E / 1786     FD8F / 1787     FD90 / 1788
     FD91 / 1789     FD92 / 178A     FD93 / 178B     FD94 / 178C
     FDB5 / 178D     FDB6 / 178E     FDB7 / 178F     FDB8 / 1790
     FDB9 / 1791     FDBA / 1792     FDBB / 1793     FDBC / 1794
     FDBD / 1795     FDBE / 1796     FDBF / 1797     FDC0 / 1798
     FDC1 / 1799     FDC2 / 179A     FDC3 / 179B     FDC4 / 179C
     FDC5 / 179D     FDC6 / 179E     FDC7 / 179F     FDC8 / 17A0
     FDC9 / 17A1     FDCA / 17A2     FDCB / 17A3     FDCC / 17A4
     FDCD / 17A5     FDCE / 17A6     FDCF / 17A7     FDD0 / 17A8
     FDD1 / 17A9     FDD2 / 17AA     FDD3 / 17AB     FDD4 / 17AC
     FE75 / 17AD     FE76 / 17AE     FE77 / 17AF     FE78 / 17B0
     FE79 / 17B1     FE7A / 17B2     FE7B / 17B3     FE7C / 17B4
     FE7D / 17B5     FE7E / 17B6     FE7F / 17B7     FE80 / 17B8
     FE81 / 17B9     FE82 / 17BA     FE83 / 17BB     FE84 / 17BC
     FE85 / 17BD     FE86 / 17BE     FE87 / 17BF     FE88 / 17C0
     FE89 / 17C1     FE8A / 17C2     FE8B / 17C3     FE8C / 17C4
     FE8D / 17C5     FE8E / 17C6     FE8F / 17C7     FE90 / 17C8
     FE91 / 17C9     FE92 / 17CA     FE93 / 17CB     FE94 / 17CC
```

- Checking the status by using SE windows by completing the following steps:

a. From the HMC, select the CPC (under Systems Management) where the CHPID or PCHID that you want to verify is, and click **Single Object Operations** (under Recovery task options).

b. On the SE, select the same CPC and click **Channels**, as shown in Figure 9-10.

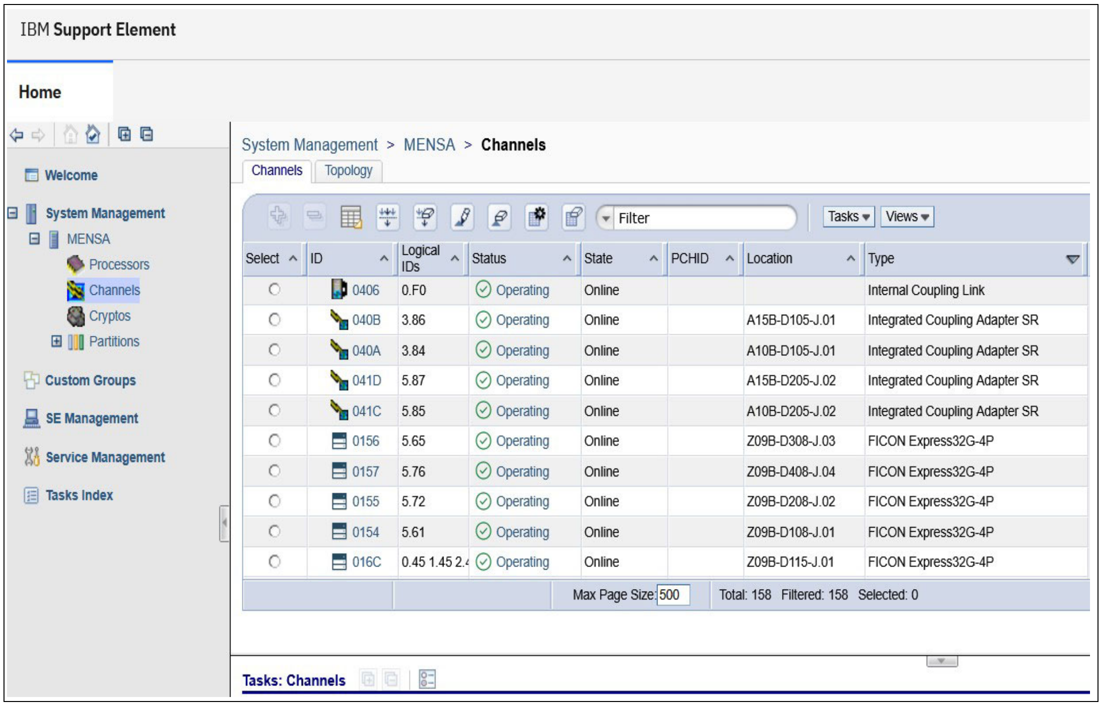

Figure 9-10 CS5 on SE Systems Management and channels

c. Look for the PCHID that you are interested in checking the status for. The result looks similar to what is shown in Figure 9-11.

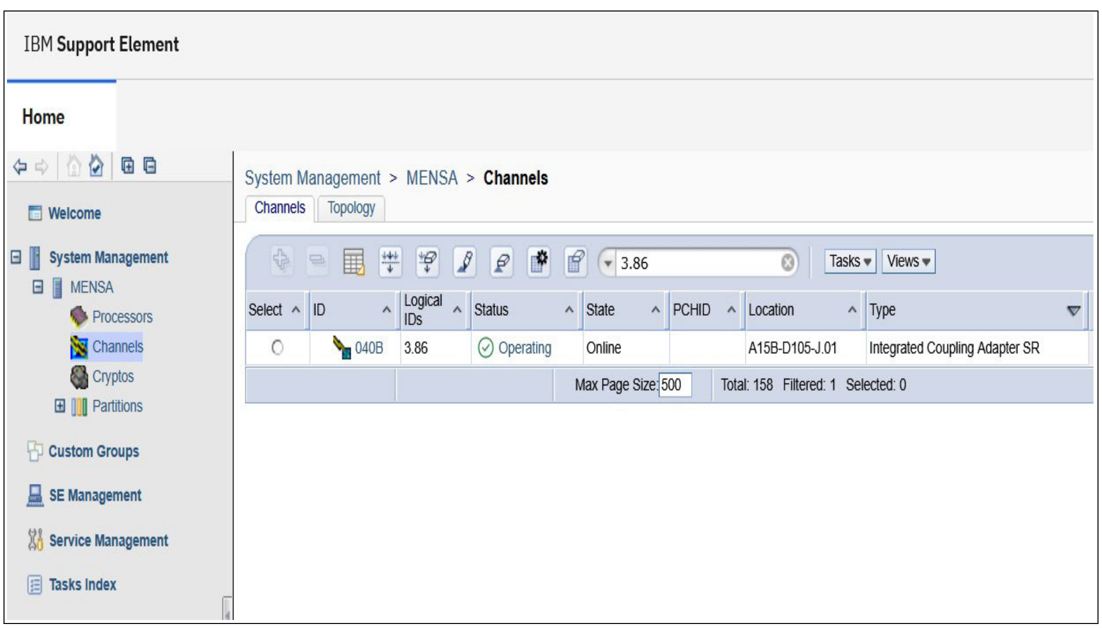

Figure 9-11 Verifying channel CS5 by using the CPC view

d. For more information about the PCHID, click the PCHID to show a result like what is shown in Figure 9-12.

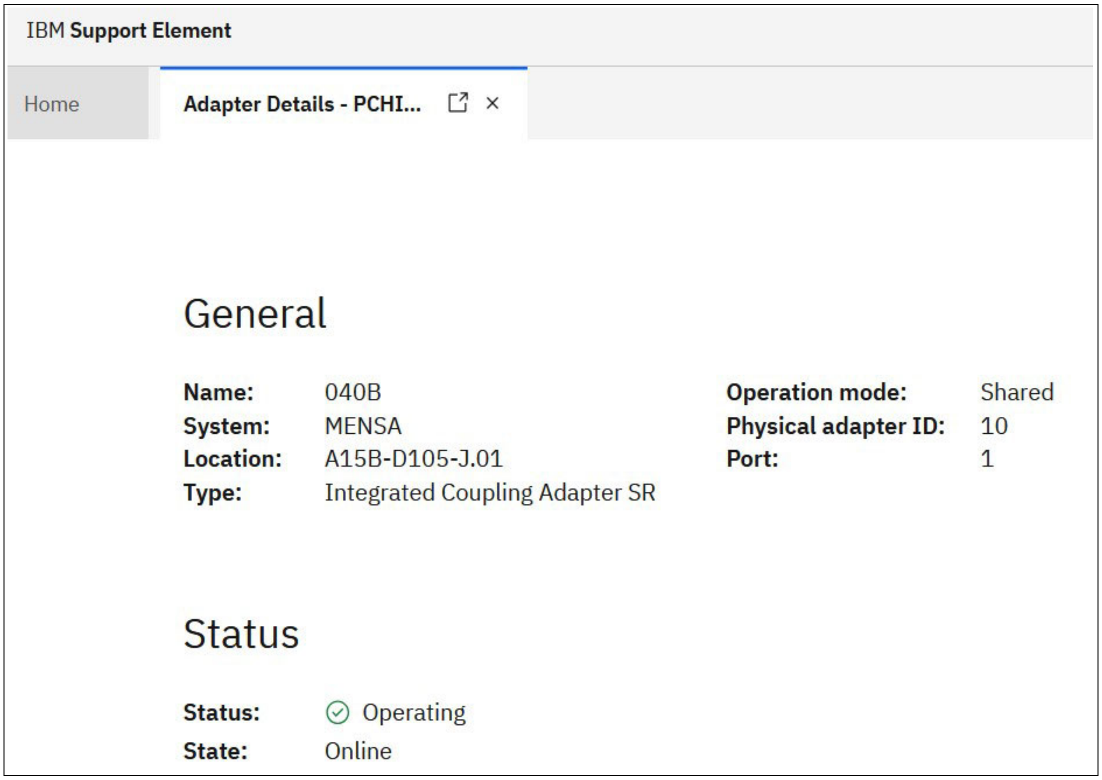

Figure 9-12 CS5 PCHID details

Another option is to check the status from the LPAR view. Select the LPAR that you are interested in checking the CHPID status of, and select the **Channels** option under that LPAR. Now, you can look for the CHPID and check the status, as shown in Figure 9-13.

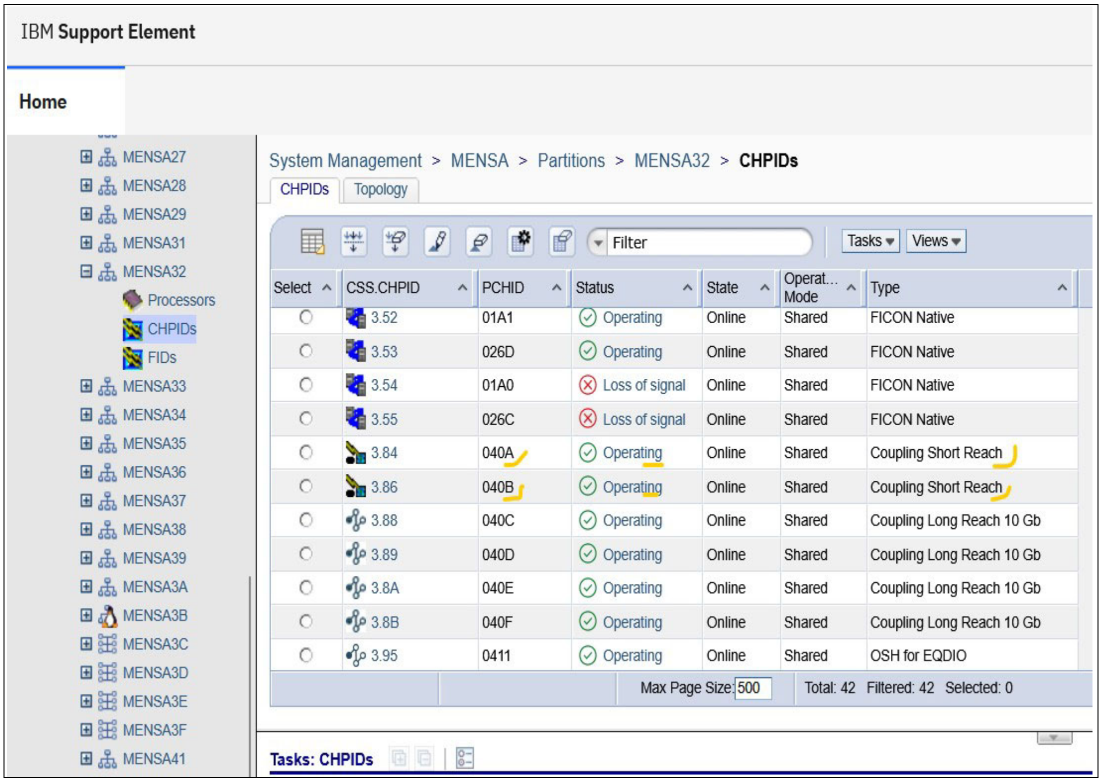

Figure 9-13 CS5 SE Verify channel LPAR view

9.4 Defining an STP timing-only link by using ICA SR

This section describes how to configure a timing-only link (for STP messages) over ICA SR connectivity.

9.4.1 STP timing-only link: Implementation

In this section, a CHPID type (CS5) connection is used to show an example of an STP timing-only link definition. The definition of the STP timing-only link is part of the *Define CF/STP link* activity that is shown in Figure 1-3 on page 5.

A coupling link connection between an IBM z17 and an IBM z16 by using CS5 CHPIDs is used as an example of how to define an STP timing-only link (see Figure 9-14).

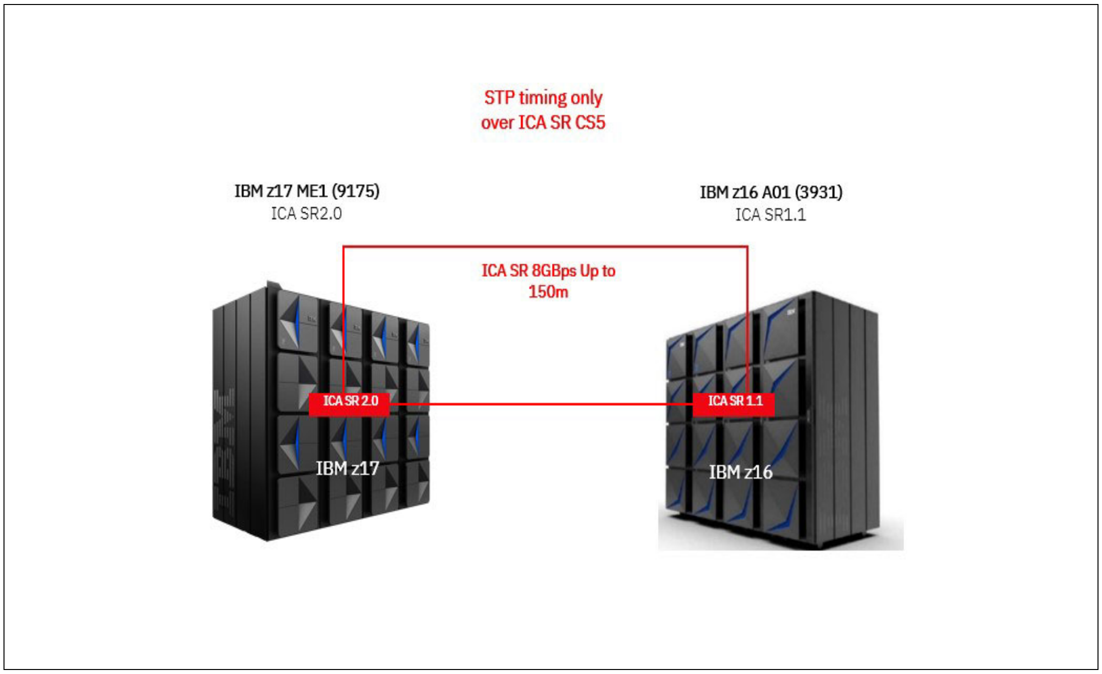

Figure 9-14 STP timing-only link connection from an IBM z17 to an IBM z16 (over CS5)

A timing-only link ICA SR is defined in the IOCDS like an ICA SR. The only difference is the CU type, which is STP in this case.

Example 9-9 shows a sample IOCP defining the CS5 CHPID for timing-only links on the IBM z17 side.

Example 9-9 IOCP defining STP timing-only link on an IBM z17 by using CS5

```
ID      MSG1='IODF79',MSG2='SYS9.IODF79 - 2025-04-17 14:34',     *
        SYSTEM=(9175,1),LSYSTEM=MENSA,                           *
        TOK=('MENSA',008003009FB89175143413700125107F00000000,00*
        000000,'25-04-17','14:34:13','SYS9','IODF79')
RESOURCE PARTITION=((CSS(0),(MENSA0A,A),(MENSA0B,B),(MENSA0C,C*
        ),(MENSA0D,D),(MENSA01,1),(MENSA02,2),(MENSA03,3),(MENSA*
        ....
        F,F),(MENSA31,1),(MENSA32,2),(MENSA33,3),(MENSA34,4),(ME*
        NSA35,5),(MENSA36,6),(MENSA37,7),(MENSA38,8),(MENSA39,9)*
        ....
CHPID PATH=(CSS(3),84),SHARED,                                   *
        PARTITION=((MENSA3E,MENSA32),(=)),CPATH=(CSS(3),85),     *
        CSYSTEM=PAVO,AID=04,PORT=1,TYPE=CS5
CHPID PATH=(CSS(3),86),SHARED,                                   *
        PARTITION=((MENSA3F,MENSA32),(=)),CPATH=(CSS(3),81),     *
        CSYSTEM=PAVO,AID=10,PORT=1,TYPE=CS5
        ....
CNTLUNIT CUNUMBR=FFFC,PATH=((CSS(3),84,86)),UNIT=STP
```

The sequence of steps to define the STP timing-only link connection between an IBM z17 and an IBM z16 by using CS5 CHPIDs through HCD is the same as for defining the CF links to a regular CF connection.

For more information about how to define CHPID Type CS5 in HCD, see 14.2.3, “Defining a Coupling Facility link with CS5 CHPIDs” on page 309.

Complete the following steps:

1. On the CF links connection step, after including the data that is related to the CPC MENSA side of the connection, type `YES` on the `Timing-only link` option of the Connect to CF Channel Path panel (see Example 9-10), and press Enter.

Example 9-10 STP timing only link from z17 to z16

```
Goto  Filter  Backup  Query  Help
 +---------------- Connect to CF Channel Path -----------------+ ---------
 |                                                             |
 |                                                             | ==> PAGE
 | Specify the following values.                               |
 |                                                             |
 | Source processor ID  . . . . . : MENSA                      |
 | Source channel subsystem ID  . : 3                          |
 | Source channel path ID . . . . : 84                         |
 | Source channel path type . . . : CS5                        |
 |                                                             |
 | Destination processor ID . . . . . . PAVO  +                |  -CU-  -#-
 | Destination channel subsystem ID . . 2   +                  |  Type  Dev
 | Destination channel path ID  . . . . E4  +                  |
 |                                                             |  CFP   8
 | Timing-only link . . . . . . . . . . YES                    |  CFP   32
 |                                                             |  CFP   8
 |                                                             |  STP
 |  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset   |  CFP   32
 |  F9=Swap    F12=Cancel                                      |  STP
 +-------------------------------------------------------------+  CFP   32
 F8        Y  ICP  SHR  N     PAVO.2    F9        Y  ICP  SHR     CFP   7
```

2. After accepting or overriding the CU and Device numbers for both processors, HCD returns to the CF Channel Path Connectivity List panel (see Example 9-11). You can see that the STP timing-only links are now connected.

Example 9-11 STP timing-only links connected

```
Goto  Filter  Backup  Query  Help
 ------------------------------------------------------------------------
                     CF Channel Path Connectivity List             Row
 Command ===> ____________________________________________ Scroll ===> P

 Select one or more channel paths, then press Enter.

 Source processor ID . . . . . : MENSA     Mensa
 Source channel subsystem ID . : 3
 Source partition name . . . . : *

   ----------Source----------   -----------Destination-----------   -CU-
 / CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type
 _ 84  04/1  N  CS5  SHR  N     PAVO.3     85  11/2  N  CS5  SHR    STP
 _ 86  10/1  N  CS5  SHR  N     PAVO.3     81  05/2  N  CS5  SHR    STP
```

9.4.2 STP timing-only links: Verifying the configuration

After activating the new configuration with the ICA SR links defined as STP timing-only links and all cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the SE windows.

The same process that you used for ICA SR links should be followed to check the status (online and operating) of the CS5 CHPIDs that were defined to work as STP timing-only links.

Complete the steps that are described in 9.3.1, “ICA SR: Verifying the configuration” on page 180 by using the respective CHPID/PCHID of the links for which you are interested in checking the status.

9.5 CF LPAR setup and Coupling Facility Control Code Level 26

The intention of this section is to remind you of the configuration aspects that are related to a CF LPAR, such as the possible sysplex configuration with an IBM z17, memory considerations due to the new CF Level, and the changes in Coupling Facility Control Code (CFCC) Level 26.

As shown in Figure 1-3 on page 5, the following aspects must be considered for the activity *Setup CF LPAR*:

- IBM z17 servers support active participation in the same Parallel Sysplex, and Coupling Link connections with, these servers:
  - IBM z16 A01 or IBM z16 A02
  - IBM z15 T01 or IBM z15 T02
  - This is both a direct CEC-to-CEC connectivity statement, as well as a presence anywhere in the sysplex regardless of direct or indirect connectivity statement:
    - Configurations with z/OS on one of these servers can add an IBM z17 server to their sysplex for either a z/OS or a CF image;
    - Configurations with a Coupling Facility on one of these servers can add a IBM z17 server to their Sysplex for either a z/OS or a Coupling Facility image.
    - Usage of Integrated Coupling Adapter Short Reach (ICA SR) or CE LR coupling links is required.
- Memory considerations:
  - Memory planning must consider the CFCC memory and structure size increases that are associated with a new level of the CFCC.
  - LPARS running the CFCC code might increase storage requirements when moving from CF Level 25 (or earlier) to CF Level 26. In fact, structure size requirements increases as the coupling facility’s internal data structures for CF structure objects must increase in size to accommodate new functions and capabilities.
  - As a best practice, CF structures should always be re-sized on any CFLEVEL upgrades. Make use of:

- The z/OSMF Sysplex Management Application, CFRM Policy Editor, CF structure sizing capability for CFLEVEL 26;
- Use the Coupling Facility Structure Sizer (CFSizer) Tool or the batch SIZER utility.

Note: For resizing your CF structures as needed, make the corresponding changes to your CFRM policy INITSIZE and SIZE values.

9.5.1 Coupling Facility Control Code Level 26

The new CFCC Level 26 introduces changes and improvements in the following areas:

- Simplification:
  - Removal of support for CF Flash Memory (VFM) and CF images using Dedicated GP processors.
- Parallel Sysplex Scalability, Virtualization, Consolidation, Density:
  - GA1 Base HW/FW enhancements.
- Increase ICP Buffers per CHPID from 7 to 8:
  - IBM z17 will always use 8 buffers;
  - Improve capacity/throughput for internal coupling channels;
  - HCD/IOCP update to allow 7 or 8 subchannels/devices per ICP CHPID (default to 8).
- Maximum of 384 coupling CHPIDs (of all types) per CEC, 64 max ICPs CHPIDs; 4 CHPIDs per port.

For more information about the enhancements that were made in CF level 25, see: *IBM z17 (9175) Technical Guide*, SG24-8579.

9.6 Dynamic I/O for Stand-alone Coupling Facility

Before IBM z14 GA2 (Driver 36), a SACF could not change the I/O configuration dynamically because of a missing HCD running in the IBM Z server. Changing the I/O configuration for such a server is cumbersome, error-prone, and requires a Power on Reset (POR).

Dynamic configuration capabilities for SACFs were introduced on IBM z14 GA2, and they are enabled by default on IBM z15, IBM z16 and IBM z17.

With this support, an activation service (firmware function) is started on the SACF CPC to perform this role, which provides simple dynamically activated I/O changes with no IML requirement.

With IBM z15, z16, and z17 and there are new firmware communication pathways from the “driving” HCD managing the input/output definition file (IODF) changes through the HMC/SE to the SACF CPC, and ultimately the Master Control Service (MCS) activation service (firmware function):

- For passing the modified target configuration
- For driving the Dynamic I/O activate and associated recovery/management functions

For more information, see *IBM z17 (9175) Technical Guide*, SG24-8579.

The dynamic activation is like the existing remote dynamic activation on a server that supports z/OS LPARs, but for hardware only changes.

For more information about how to implement this new support, see 5.5.2, “Remote dynamic activation of I/O configurations for Stand Alone Coupling Facilities (SACFs), Linux on IBM Z, and z/TPF” on page 117.

9.6.1 IBM Processor Resource/System Manager solution

With IBM z15, IBM z16 and IBM z17, dynamic activation of a new or changed IODF on a SACF CPC is supported:

- Without requiring a POR/IML
- Without requiring the presence (on the same CPC) of any z/OS or z/VM image running an HCD instance

This is a base IBM Processor Resource/System Manager (PR/SM) solution; it does not require the usage of Dynamic Partition Manager (DPM) mode:

- A firmware function is used, which is a firmware-based appliance version of the HCD instance.
- The solution is fully managed by the IBM Z firmware.
- The solution is included with the base firmware (no need to order a feature code).

The firmware LPAR is activated by default on IBM z15, IBM z16 and IBM z17.

The firmware appliance LPAR on a supported system is driven by an updated HCD running in z/OS LPAR on a remote IBM Z, as shown on Figure 9-15.

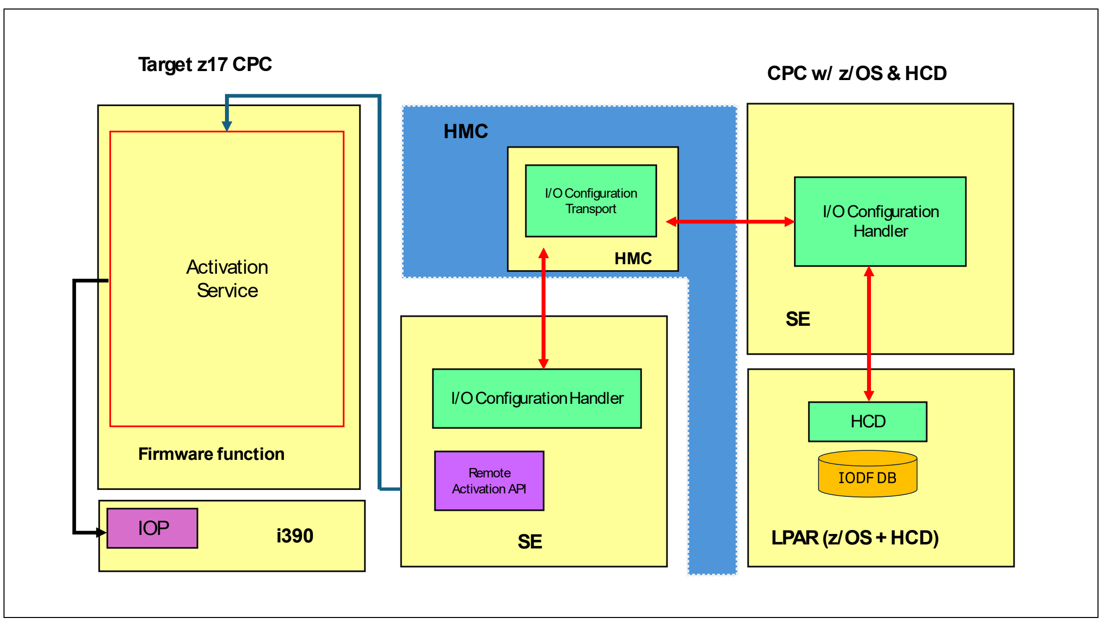

Figure 9-15 Dynamic I/O for a Stand-alone Coupling Facility

10

Specialized features

This chapter covers the configuration of all specialized features that are available for IBM z17 systems.

> **Note:** The examples that are shown in this chapter are based on the IBM z17 ME1 (9175). However, these examples can also be used with the IBM z16 A01 (3931), the IBM z16 A02 and the IBM z16 AGZ (3932).

This chapter includes the following topics:

- Crypto Express8S
- Virtual Flash Memory

10.1 Crypto Express8S

This section provides information about the configuration of the Crypto Express8S feature on IBM z17. It covers cryptographic domains, configuration rules, and what to consider when you are planning for a nondisruptive installation of cryptographic features.

The chapter includes the steps for defining and configuring the Crypto Express8S feature to a logical partition (LPAR).

This section includes the following topics:

- Crypto Express8S overview
- Planning for a Crypto Express8S configuration
- Configuring Crypto Express8S
- Handling cryptographic coprocessors by using ICSF

10.1.1 Crypto Express8S overview

There are two generations of cryptographic coprocessors[1] that are supported for IBM z17:

- Crypto Express7S (Feature Code 0899 (one adapter on card) and Feature Code 0898 (two adapters on card)), carry forward only (MES from IBM z15 or IBM z16)
- Crypto Express8S (Feature Code 0909 (one adapter on card) and Feature Code 0908 (two adapters on card)), new built or carry forward (MES from IBM z16)

This section describes the Crypto Express8S features for the IBM z17 system.

Each cryptographic coprocessor of an IBM z17 ME1 has 85 physical sets of registers, which correspond to the maximum number of LPARs running on an IBM z17 ME1. Each of these sets belongs to a domain as follows:

- A cryptographic domain index, in the range of 0 - 84 is allocated to an LPAR by the definition of the partition in its image profile. The same domain must also be allocated to the Integrated Cryptographic Service Facility (ICSF) instance running in the LPAR that uses the Options data set.
- Each ICSF instance accesses only the master keys that correspond to the domain number that is specified in the LPAR image profile at the Support Element (SE) and in its Options data set. Each ICSF instance sees a logical cryptographic coprocessor that consists of the physical cryptographic engine and the unique set of registers (the domain) that is allocated to this LPAR.

[1] Regional Crypto Enablement (RCE) is also supported, but it is not covered here.

The installation of CP Assist for Cryptographic Functions (CPACF) Enablement (Feature Code 3863) is one of the prerequisites for using the special hardware cryptographic feature in an IBM z17. Feature Code 3863 enables the following functions:

- For data privacy and confidentiality: DES, TDES, and Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys. Also with Feature Code 3863, the Modulo arithmetic unit is enabled, which supports Elliptic Curve Cryptography for verification and signing by using NIST P256, P384, and P521 curves; Brainpool Curves P256, P384, and P512; and Edward curves Ed25519 and Ed448 with clear keys.
- For data integrity: Secure Hash Algorithm-1 (SHA-1) 160-bit, and SHA-2 for 224-, 256-, 384-, and 512-bit support. SHA-3 for 224-, 256- 384-, and 512-bit support, and SHAKE for 128- and 256-bit support. SHA-1, SHA-2, and SHA3 are enabled on all IBM z17 ME1 systems, and they do not require the no-charge enablement Feature Code 3863.
- For Key Generation: Pseudo-Random Number Generation (PRNG), Deterministic Random Number Generation (DRNG), and True Random Number Generation (TRNG).
- For message authentication code: Single-key message authentication code and double-key message authentication code.

The total number of cryptographic Peripheral Component Interconnect Express (PCIe) adapters (combined Crypto Express8S or Crypto Express7S features) cannot exceed 60 per IBM z17. Up to 16 features with one HSM adapter can be used (Crypto Express7S or Crypto Express8S), or up to 30 features with two HSM adapters (Crypto Express7S 2 port or Crypto Express8S 2 port) resulting in 60 HSM adapters.

The initial order for Crypto Express8S is two features (two HSM adapters for Feature Code 0909, and four HSM adapters for Feature Code 0908). After the initial order, the minimum order is one feature.

Each Crypto Express8S (Feature Code 0909) contains one HSM adapter, and each Crypto Express8S (Feature Code 0908) contains two HSM adapters. The adapter can be in the following configurations:

- Common Cryptographic Architecture (CCA) Coprocessor (CEX8C)
- Public Key Cryptography Standards (PKCS) #11 (EP11) Coprocessor (CEX8P)
- Accelerator (CEX8A)

During the feature installation, the PCIe adapter is configured by default as the CCA coprocessor.

The configuration of the Crypto Express8S adapter as an EP11 coprocessor requires a Trusted Key Entry (TKE) workstation Hardware (Feature Code 0057 for the rack-mounted workstation, and Feature Code 0058 for the tower workstation) with TKE 10.1 Licensed Internal Code (LIC) (Feature Code 0883). Older TKE FCs 0085 and 0086 TKEs can be carried forward and converted by replacing the installed Crypto adapter card with a 4770 Crypto adapter card FC 0851.

The Crypto Express8S feature does not use channel path IDs (CHPIDs) from the channel subsystem (CSS) pool. However, the Crypto Express8S feature requires one slot in a PCIe I/O drawer, and one physical channel ID (PCHID) for each PCIe cryptographic adapter.

Table 10-1 summarizes the cryptographic feature codes for IBM z17.[2]

Table 10-1 Cryptographic feature codes

Feature Code	Description
3863	CPACF enablement This feature is a prerequisite to use CPACF (except for SHA-1, SHA-2, and SHA-3) and cryptographic coprocessor hardware.
0908	Crypto Express8S (2-port) card A maximum of 30 features for an IBM z17 ME1 can be ordered (a minimum of two adapters). This feature is optional, and each feature contains two PCIe cryptographic adapters (an adjunct processor). This feature is supported by IBM z16 and IBM z17.
0909	Crypto Express8S (1-port) card A maximum of 16 features can be ordered (a minimum of two adapters). This feature is optional, and each feature contains one PCIe cryptographic adapter (an adjunct processor). This feature is supported by only IBM z16 and IBM z17.
0898	Crypto Express7S (2-port) card This feature cannot be ordered for a new IBM z17 ME1 system, but only on a carry-forward MES from an IBM z15 or an IBM z16 system. A maximum of 30 features can be carried forward for an IBM z17 ME1(a minimum of two adapters for all models). This feature is optional, and each feature contains two PCIe cryptographic adapters (an adjunct processor). This feature is supported by IBM z17, IBM z16 and IBM z15.
0899	Crypto Express7S (1-port) card This feature cannot be ordered for a new IBM z17 ME1 system, but only on a carry-forward MES from an IBM z15 or an IBM z16 system. A maximum of 16 features can be carried forward (a minimum of two adapters). This feature is optional, and each feature contains one PCIe cryptographic adapter (an adjunct processor). This feature is supported by only IBM z17, IBM z16 and IBM z15.
0058	TKE tower workstations A TKE provides basic key management (key identification, exchange, separation, updates, and backup) and security administration. It is optional for running a Crypto Express feature in CCA mode in a non-PCI-compliant environment. It is required for running in EP11 mode and CCA mode with full PCI compliance. The TKE workstation has one Ethernet port, and supports connectivity to an Ethernet local area network (LAN) operating at 10, 100, or 1000 Mbps. It also requires Feature Code 0157. Up to 10 features that are combined with Feature Code 0057 per IBM z17 can be ordered.
0157	TKE Table Top Keyboard/Monitor/Mouse A tabletop monitor with a US English language keyboard. There is a touchpad for pointing, and a country-specific power cord.
0057	TKE rack-mounted workstation The rack-mounted version of the TKE, which needs a customer-provided standard 19-inch rack. It also requires Feature Code 0156. When using smart card readers, an extra customer-provided tray is needed. Up to 10 features combined with Feature Code 0058 per IBM z17 can be ordered.
0156	TKE Rack Keyboard/Monitor/Mouse A 1U rack-mounted display and keyboard with a built-in pointing device. The keyboard comes in the US English language.

[2] Other IBM Z servers are also listed where applicable.

Feature Code	Description
0883	TKE 10.1 LIC Included with the TKE tower workstation Feature Code 0058 and the TKE rack-mounted workstation Feature Code 0057 for IBM z17. Earlier versions of TKE features (Feature Codes 0085 and 0086) can also be upgraded to the TKE 10.1 LIC by replacing the installed Crypto adapter card with a 4770 Crypto adapter card FC 0851 if the TKE is assigned to a IBM z15 or later system.
0886	TKE Smart Card Reader Access to information in the smart card is protected by a PIN. One Feature Code includes two smart card readers, two cables to connect to the TKE workstation, and 20 smart cards. Older Smart Card Readers FC 0885 or 0891 can be carried forward.
0889	TKE extra smart cards When one Feature Code is ordered, 10 smart cards are included. The order increment is 1 - 99 (990 blank smart cards). Older Smart Cards FC 0892 and 0900 can be carried forward.

Note: You might need the TKE workstation including TKE Smart Card Reader while you run CEX in CCA mode to meet certain security standards requirements.

For more information about the Crypto Express8S feature and the corresponding crypto features, see: *IBM z17 (9175) Technical Guide*, SG24-8579.

10.1.2 Planning for a Crypto Express8S configuration

Note: New ICSF support is required to administer a Crypto Express8S coprocessor that uses a TKE workstation, due to leveraging quantum algorithms. Otherwise, existing workloads run on IBM z17 ME1 without requiring ICSF support.

Leveraging of the new function is supplied in z/OS 2.5 or z/OS 3.1 (base, which is HCR77D2). When leveraging new quantum-safe algorithms and sharing a KDS in a sysplex, ensure that HCR77D2 is installed on all systems.

All supported levels of ICSF automatically detect what hardware cryptographic capabilities are available where it is running, then enables functions as needed. No toleration of new hardware is necessary. If you want to leverage new capabilities, then ICSF support is necessary.

For the latest MCL bundle requirements, see the *Driver Level 61 Exception Letter*.

IBM z17 ME1 always operates in LPAR mode. The concept of *dedicated coprocessor* does not apply to the PCIe adapter. A PCIe adapter, whether configured as a coprocessor or accelerator, is made available to LPARs as directed by the domain assignment and the candidate list. This process occurs regardless of the shared or dedicated status that is given to the central processors (CPs) in the partition.

IBM z17 ME1 enables up to 85 LPARs to be active concurrently. Each PCIe adapter on a Crypto Express8S feature supports 85 domains, whether it is configured as a Crypto Express8S coprocessor or a Crypto Express8S accelerator.

For availability reasons, the minimum configuration consists of two Crypto Express8S features so that every potential LPAR can have access to two cryptographic adapters on two different cards.

More Crypto Express8S features might be needed to satisfy application performance and availability requirements:

- For availability, spread the assignment of multiple PCIe adapters of the same type (accelerator or coprocessor) to one LPAR across features in multiple I/O domains.
- Using retained private keys on a PCIe adapter that is configured as a Crypto Express8S coprocessor creates an application single point of failure. This point of failure exists because RSA-retained private keys are not copied or backed up.
- There is an intrusion latch within the PCIe adapter logic that is set when the feature is removed from the system. If the feature is reinstalled and power is applied, the coprocessor keys and secrets are zeroed and the intrusion latch is reset.

  If a TKE workstation is available, disable the PCIe adapter before you remove the feature from the system. When the feature is reinstalled, the coprocessor keys and secrets are not zeroed. The intrusion latch is reset, and the coprocessor remains in the disabled state. Then, the PCIe adapter can be enabled from the TKE and normal operations can resume.

  Plan the definition of domain indexes and cryptographic coprocessor numbers in the candidate list for each LPAR to prepare the cryptographic configuration. You can also define or change that cryptographic definition dynamically to an active LPAR with a running system. For more information, see “Change LPAR Cryptographic Controls function” on page 201.
- Crypto Express8S features can be installed concurrently when all physical requirements are fulfilled. Dynamically enabling a new PCIe adapter to a partition requires these configurations:
  - At least one usage domain index must be defined to the LPAR.
  - The cryptographic coprocessor numbers must be defined in the partition candidate list.
- The same usage domain index can be defined more than once across multiple LPARs. However, the cryptographic coprocessor number that is coupled with the usage domain index that is specified must be unique across all *active* LPARs.

  The same cryptographic coprocessor number and usage domain index combination can be defined for more than one LPAR. This feature can be used, for example, to define a configuration for backup situations. In this case, only one of the LPARs can be active at any one time.
- Newly installed Crypto Express8S features are assigned coprocessor numbers sequentially during the Power on Reset (POR) that follows the installation.

  However, when a Crypto Express8S feature is installed concurrently by using the Nondisruptive Hardware Change task, the installation might select an out-of-sequence coprocessor number from the unused range. In this case, communicate the cryptographic coprocessor numbers that you want to use to the IBM installation team.

  When the task is used to concurrently remove a PCI cryptographic feature, the coprocessor number is automatically freed.

Table 10-2 is a simplified configuration map for IBM z17 ME1. Each row identifies a PCIe adapter, and each column identifies a domain index number. Each cell entry indicates the LPAR to be assigned to the cryptographic coprocessor number that is coupled with the usage domain index.

Table 10-2 Planning for logical partitions, domains, and PCIe adapter numbers

PCIe adapter	Domain index 0	Domain index 1	Domain index 2	.../...	Domain index 84
PCIe adapter 0	LP00 LP02	LP04	LP05		
PCIe adapter 1	LP01 LP02				
PCIe adapter 2					
…					
…					
PCIe adapter 57					
PCIe adapter 58					
PCIe adapter 59					

Table 10-2 illustrates these characteristics:

- LPARs LP00 and LP01 use domain 0 (zero), but are assigned different PCIe adapters. There is no conflict. They can be concurrently active.
- LPAR LP02 uses domain 0 (zero) on the set of cryptographic adapters that are already defined to LP00 and LP01. Therefore, LP02 cannot be active concurrently with either LP00 or LP01. However, the definition might be valid for backup situations.
- LPARs LP04 and LP05 use different domain numbers for PCIe cryptographic adapter 0 (zero), so there is no conflict. The combination of domain number and cryptographic coprocessor number is unique across partitions.

> **Important:** Any combination of PCIe adapter and domain index should contain only one active LPAR. The combination of cryptographic coprocessor number and usage domain index must be unique across all *active* LPARs.

For more information about the Crypto Express8S feature for IBM Z, see *IBM z17 (9175) Technical Guide*, SG24-8579.

10.1.3 Configuring Crypto Express8S

This section provides steps for configuring Crypto Express8S for IBM z17 ME1.

The IBM z17 ME1 operates only in LPAR mode. For each LPAR that requires access to a PCIe adapter, which is configured as either an accelerator or a coprocessor, the required information must be defined in the partition image profile. This technique ensures the correct usage of the cryptographic features when the associated partition is activated.

Concurrent changes to the Crypto Express8S features and controls when the partition is already activated are provided by special functions on the SE.

Checking whether the CPACF Enablement feature is installed

Feature Code 3863 on the IBM z17 ME1 enables various algorithms (see page 191 for a list of supported algorithms) by the CPACF. It is one of the prerequisites for using the Crypto Express8S feature. Verify whether the CPACF feature is correctly installed on the processor before you configure cryptographic functions. This information is displayed in the SE and can be verified by completing the following steps:

1. Log on to the SE directly or click **Single Object Operations** from the Hardware Management Console (HMC).
2. Open the **System details** menu of the central processor complex (CPC). The system details window opens (Figure 10-1).

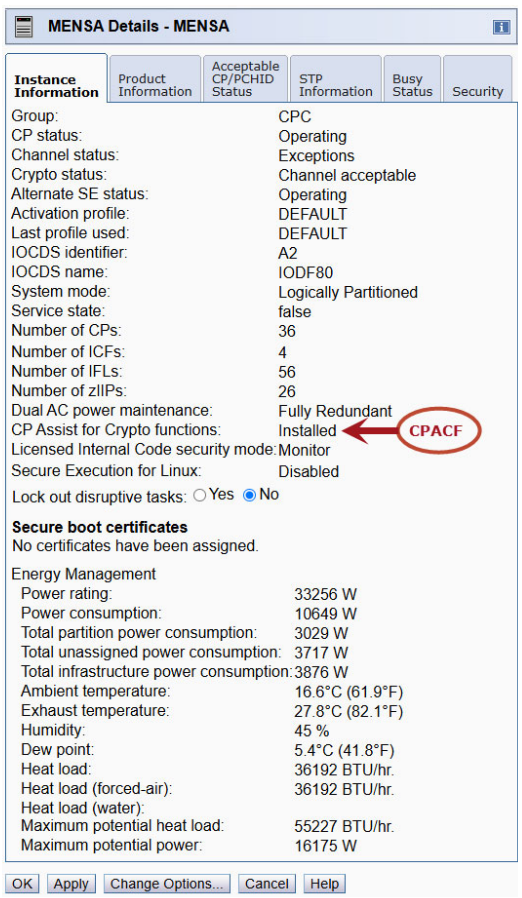

Figure 10-1 System details: CPACF installed

3. Click the **Instance Information** tab, and verify that the CPACF Enablement Feature Code 3863 is installed:
   – If the window shows the message “CP Assist for Crypto Functions: Installed”, the CPACF enablement Feature Code 3863 is enabled.
   – If the window shows the message “CP Assist for Crypto Functions: Not installed”, Feature Code 3863 is not installed. You can still customize the partition image profiles, but the cryptographic functions do not operate.

Logical partition cryptographic definition

The next step is to define the following cryptographic resources in the image profile for each LPAR:

- Usage domain index
- Control domain index
- PCI Cryptographic Coprocessor Candidate List
- PCI Cryptographic Coprocessor Online List

This task is accomplished by using the Customize/Delete Activation Profile task, which is in the Operational Customization Group, either from the HMC or from the SE. Modify the cryptographic initial definition from the Crypto option in the image profile, as shown in Figure 10-2. After this definition is modified, any change to the image profile requires you to DEACTIVATE and ACTIVATE the LPAR for the change to take effect. Therefore, this cryptographic definition is disruptive to a running system.

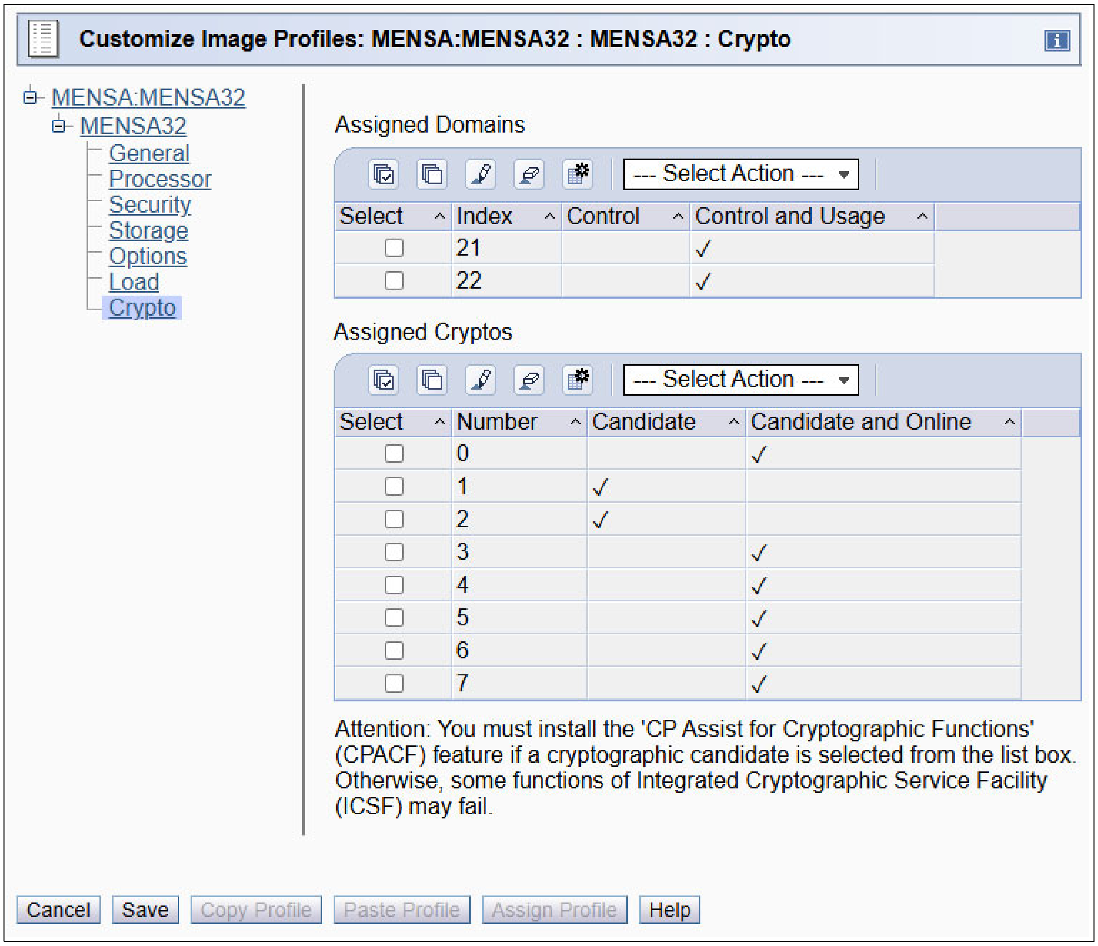

Figure 10-2 Customize Image Profiles: Crypto

> **Tip:** Operational changes can be made by using the Change LPAR Cryptographic Controls task from the SE, which reflects the cryptographic definitions in the image profile for the partition. With this function, you can dynamically add and remove the cryptographic feature without stopping a running operating system (OS). For more information about using this function, see "Change LPAR Cryptographic Controls function" on page 201.

The cryptographic resource definitions have the following meanings:

- Control Domain

Identifies the cryptographic coprocessor domains that can be administered from this LPAR if it is being set up as the TCP/IP host for the TKE.

If you are setting up the host TCP/IP in this LPAR to communicate with the TKE, the partition is used as a path to other domains' Master Keys. Indicate all the control domains that you want to access (including this partition's own control domain) from this partition.

- Control and Usage Domain

  Identifies the cryptographic coprocessor domains that are assigned to the partition for all cryptographic coprocessors that are configured on the partition. The usage domains cannot be removed if they are online.

  The numbers that are selected must match the domain numbers that are entered in the Options data set when you start this partition instance of ICSF.

  The same usage domain index can be used by multiple partitions regardless to which CSS they are defined. However, the combination of PCIe adapter number and usage domain index number must be unique across all active partitions.

- Cryptographic Candidate list

  Identifies the cryptographic coprocessor numbers that are eligible to be accessed by this LPAR. From the list, select the coprocessor numbers 0 - 15 that identify the PCIe adapters to be accessed by this partition.

  No error condition is reported when a cryptographic coprocessor number, which is selected in the partition candidate list, is available to the partition when the partition is activated, either because it is configured off or not installed. The cryptographic coprocessor number is ignored and the activation process continues.

  If the cryptographic coprocessor number and usage domain index combination for the coprocessor that is selected is already in use by another active LPAR, the activation of the LPAR fails (see Figure 10-3 on page 198). In this conflicting case, you must review the cryptographic information for all active LPARs from the **Summary** tab of the View LPAR Cryptographic Controls task (see Figure 10-5 on page 201). Resolve the error based on the collected data by assigning a unique combination of PCIe adapter number and usage domain index number.

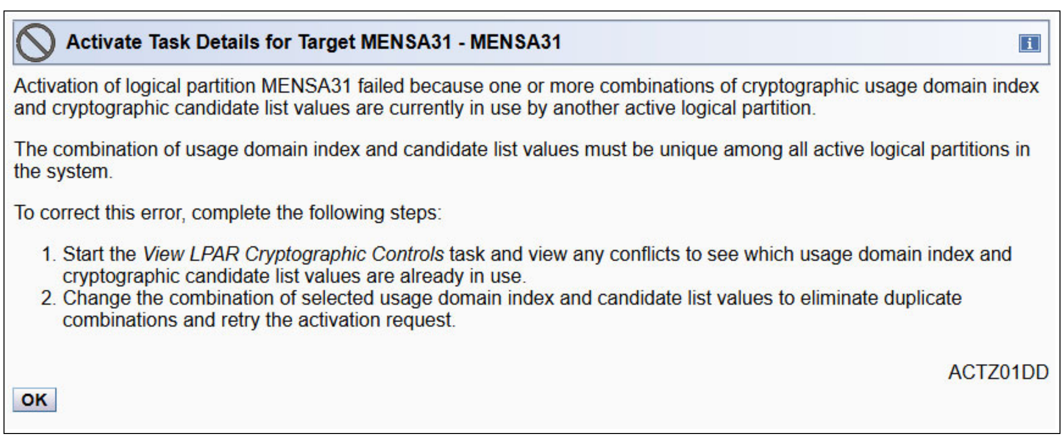

Figure 10-3 Activation of LPAR failed: ACTZ01DD

- Cryptographic Online list

  Identifies the cryptographic coprocessor numbers that are automatically brought online during LPAR activation. The numbers that are selected in the online list must also be part of the candidate list.

After the next partition activation, installed PCI Cryptographic Coprocessors that are in the partition's PCI Cryptographic Coprocessor Candidate list but not on the PCI Cryptographic Coprocessor Online list are in a *configured off* state (Standby). They can later be configured online to the partition by selecting **Configure On/Off** from the SE. For more information, see "Configuring a Crypto Express8S online or offline on a logical partition" on page 212.

When the partition is activated, no error condition is reported if a cryptographic coprocessor number that is selected in the partition's online list is not installed. The cryptographic coprocessor is ignored and the activation process continues.

When a cryptographic coprocessor number that is selected in the partition's online list was previously configured to an *off* state to the partition, it is automatically configured back to an *on* state when the partition is activated. The cryptographic online list is always selected from the image profile for each LPAR.

Cryptographic configuration by using the Support Element

From the SE, you can do these tasks:

- Display PCI Cryptographic Configuration.
- Display LPAR cryptographic controls (domain index and candidate/online lists for activated partitions).
- Reconfigure the coprocessor from and to the accelerator.
- Configure a cryptographic coprocessor and accelerator on or off to an LPAR.
- Change LPAR cryptographic controls to an LPAR.

These tasks require you to work from the SE. To get to the appropriate SE task, log on to the SE directly or click **Single Object Operations** from the HMC.

Cryptographic management

After you select the CPCs, click **Cryptographic Management** in the Configuration section.

Figure 10-4 shows the Cryptographic Management window. Use this window to obtain the installed cryptographic configuration (the association of the cryptographic number and the card serial number).

- View installed cryptographic features, including their statuses and assigned PCHIDs and coprocessor numbers. Each PCIe adapter is assigned to a coprocessor number in the range 0 - 59 as part of the configuration process. The assignment is made when the feature is installed.
- View coprocessor numbers that are still assigned to removed cryptographic features.
- Initiate the release of coprocessor numbers. Remove the relationship only when a Crypto Express8S feature is permanently removed from the CPC.

  The release option removes the relationship between a PCI cryptographic feature serial number and the assigned coprocessor numbers. Removing the relationship frees the coprocessor numbers, making them available to be assigned to a new feature serial number.

> **Important:** The coprocessor numbers are assigned to the feature serial number, *not* to the installed location. If a feature is removed from one location to be reinstalled in another location, the coprocessor number assignment remains.

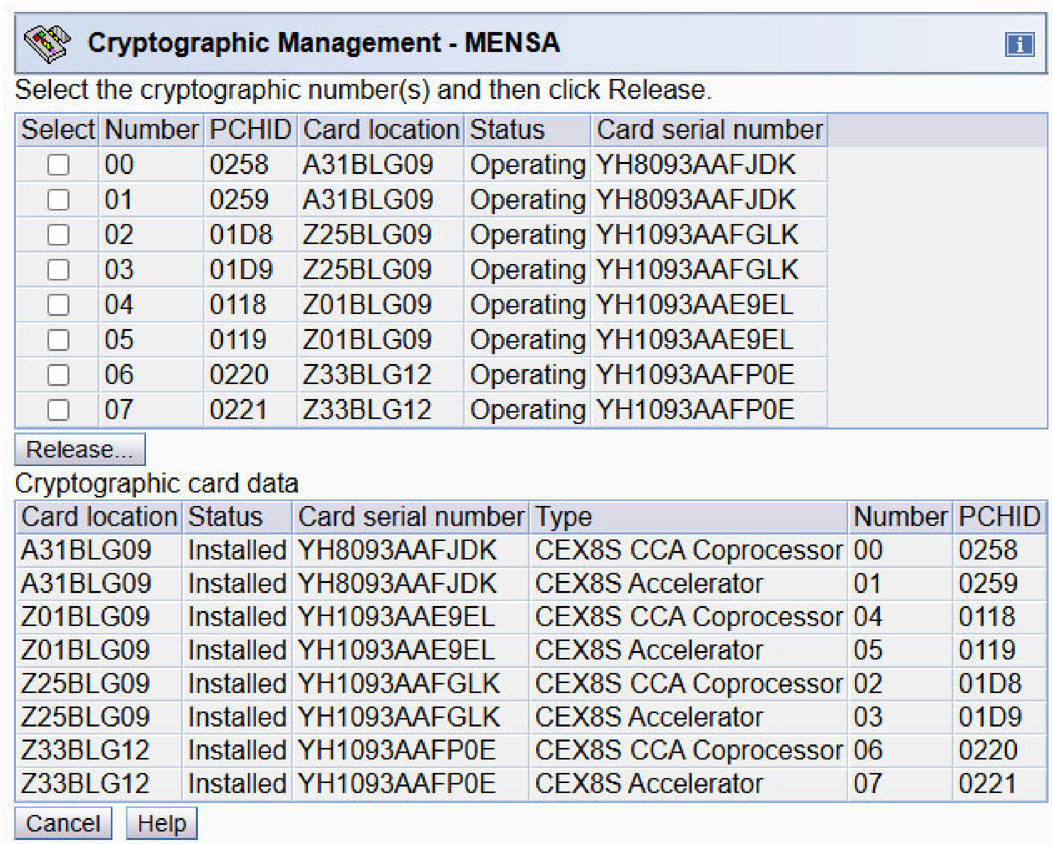

Cryptographic Management - MENSA

Select the cryptographic number(s) and then click Release.

Select	Number	PCHID	Card location	Status	Card serial number
☐	00	0258	A31BLG09	Operating	YH8093AAFJDK
☐	01	0259	A31BLG09	Operating	YH8093AAFJDK
☐	02	01D8	Z25BLG09	Operating	YH1093AAFGLK
☐	03	01D9	Z25BLG09	Operating	YH1093AAFGLK
☐	04	0118	Z01BLG09	Operating	YH1093AAE9EL
☐	05	0119	Z01BLG09	Operating	YH1093AAE9EL
☐	06	0220	Z33BLG12	Operating	YH1093AAFP0E
☐	07	0221	Z33BLG12	Operating	YH1093AAFP0E

Release...

Cryptographic card data

Card location	Status	Card serial number	Type	Number	PCHID
A31BLG09	Installed	YH8093AAFJDK	CEX8S CCA Coprocessor	00	0258
A31BLG09	Installed	YH8093AAFJDK	CEX8S Accelerator	01	0259
Z01BLG09	Installed	YH1093AAE9EL	CEX8S CCA Coprocessor	04	0118
Z01BLG09	Installed	YH1093AAE9EL	CEX8S Accelerator	05	0119
Z25BLG09	Installed	YH1093AAFGLK	CEX8S CCA Coprocessor	02	01D8
Z25BLG09	Installed	YH1093AAFGLK	CEX8S Accelerator	03	01D9
Z33BLG12	Installed	YH1093AAFP0E	CEX8S CCA Coprocessor	06	0220
Z33BLG12	Installed	YH1093AAFP0E	CEX8S Accelerator	07	0221

Cancel Help

Figure 10-4 SE Cryptographic Management

View LPAR Cryptographic Controls task

You can view active partition cryptographic definitions from the SE by selecting the CPCs and clicking **View LPAR Cryptographic Controls** in the Operational Customization pane.

The window that opens shows the definition of Usage and Control domain indexes, and PCI Cryptographic candidate and online lists. The information is provided only for active LPARs.

> **Tip:** You can review the PCI Cryptographic candidate lists and usage domain indexes that are assigned for all active LPAR from the **Summary** tab (see Figure 10-5 on page 201). The usage domain index, with the cryptographic number that is selected in the candidate list, must be unique across all partitions that are defined to the CPC. Therefore, this new tab is useful when you define or change the usage domain index for an LPAR.

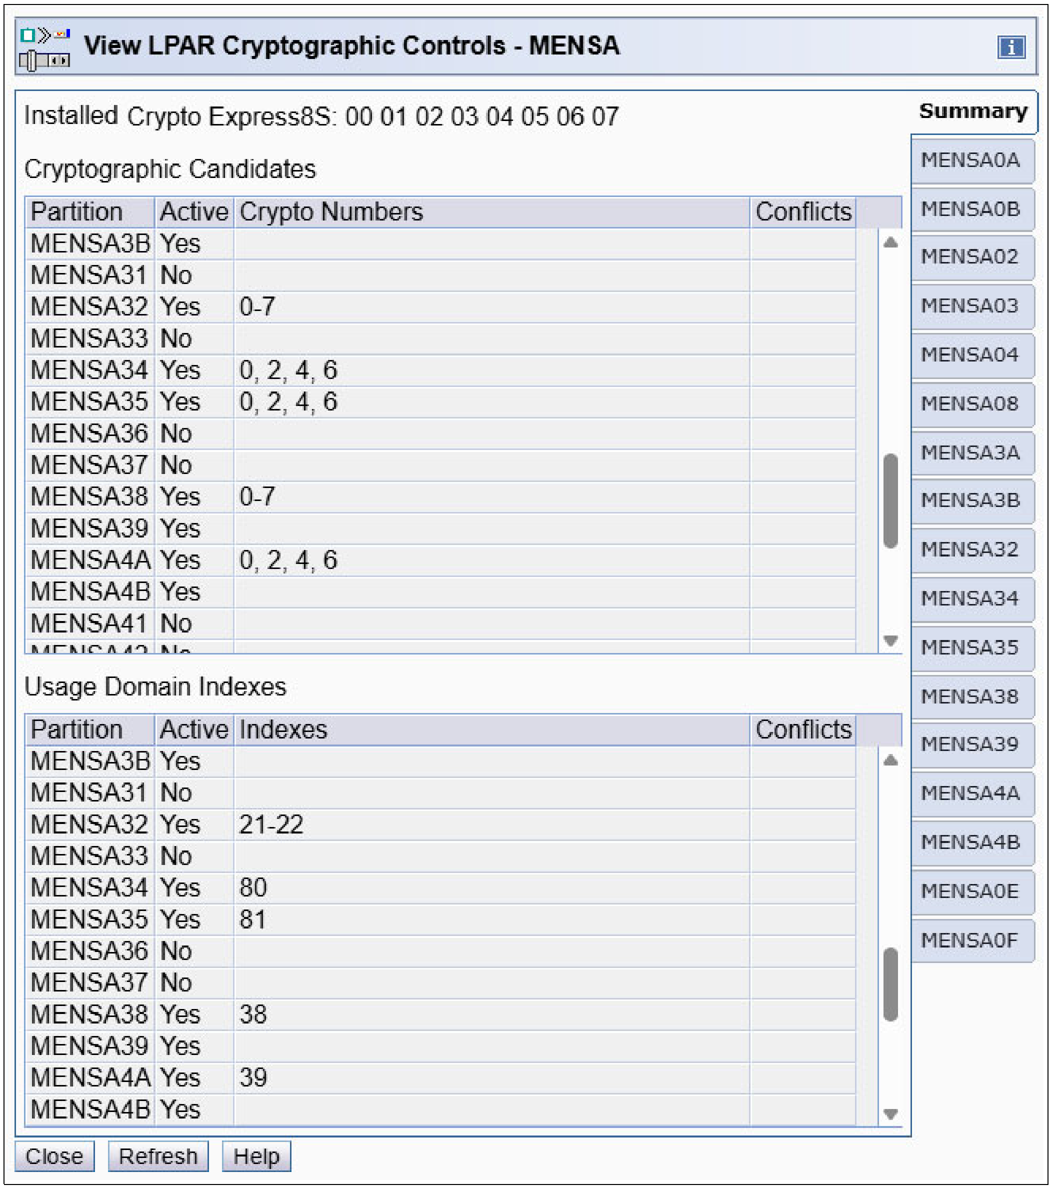

Figure 10-5 View LPAR Cryptographic Controls

This window is for informational purposes only. You can see the definitions, but you cannot change them by using this window. Modifying the cryptographic coprocessor on/off status requires the Configure On/Off task, which is described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212.

Change LPAR Cryptographic Controls function

For each LPAR, you can define these attributes:

- Usage domain index
- Control domain index
- Cryptographic Coprocessor Candidate list
- Cryptographic Coprocessor Online list

By using the Change LPAR Cryptographic Controls function, which is included in the SE for IBM z17 ME1, you can do these tasks:

- Add a cryptographic coprocessor to an LPAR for the first time.
- Add a cryptographic coprocessor to an LPAR that uses a cryptographic coprocessor.
- Remove a cryptographic coprocessor from an LPAR.
- Zeroize or clear the cryptographic secure keys for a usage domain.

Dynamic assignment of the cryptographic definition to the partition

All the cryptographic functions that are defined in the image profile can be dynamically changed by using the Change LPAR Cryptographic Controls task at the SE. For more information about defining functions in the image profile, see "Logical partition cryptographic definition" on page 196.

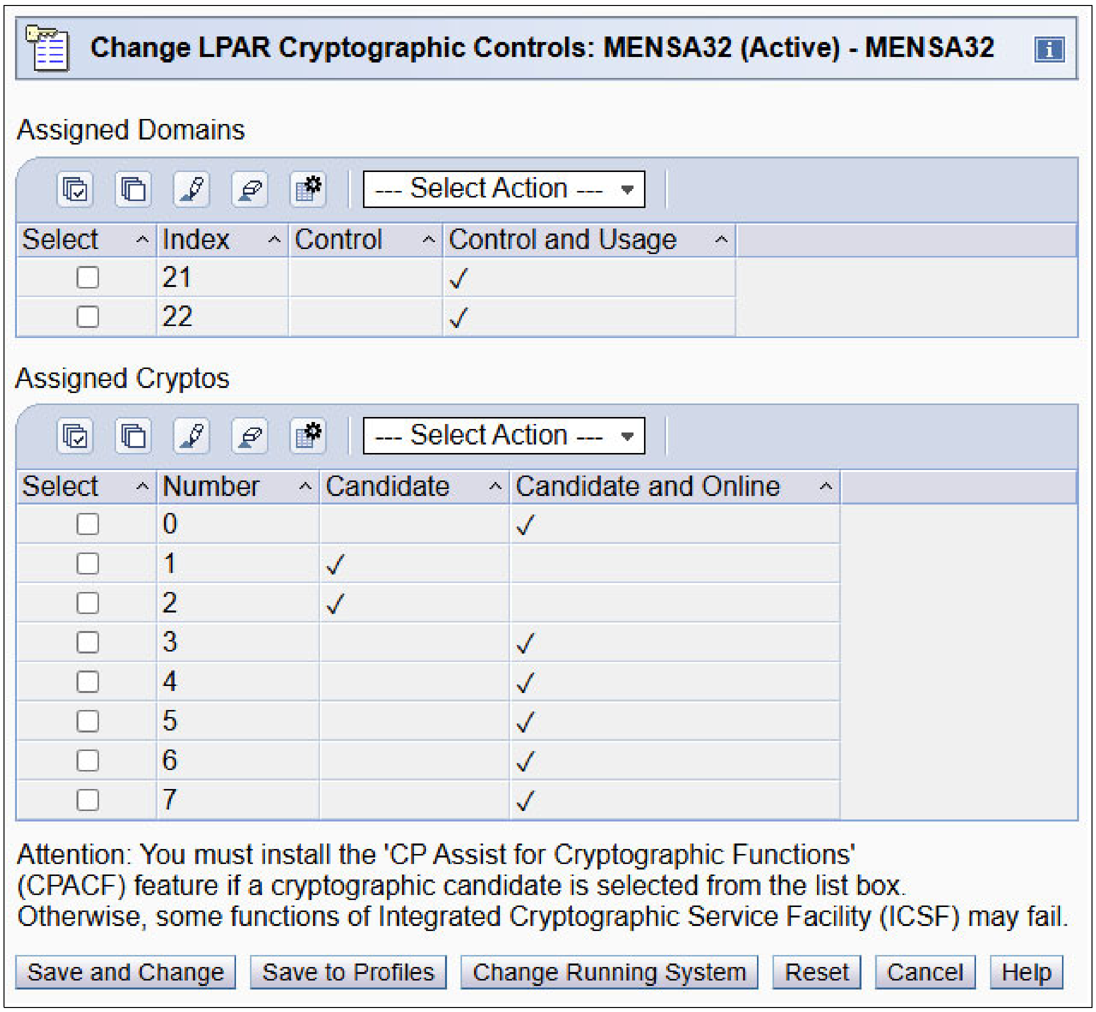

Figure 10-6 Change LPAR Cryptographic Controls: Change Running System

Select **Control** and **Usage** for each domain and **Candidate** and **Online** for each crypto (see Figure 10-6 on page 202).

After selecting the appropriate boxes, you can do these tasks:

- Save these settings to the image profile without changing the running system.
- Change the running system without saving the definition to the image profile, which means after a reactivation of the partition your changes are lost.

> **Remember:** Changes to the Cryptographic Coprocessor Online List are ignored when this option is selected.

- Save the definitions to the image profile and activate the changes immediately to the partition.

When you add or change the control or usage domain index and cryptographic coprocessor number dynamically for a running system, a confirmation message appears. After processing, a status window opens and indicates the result of a dynamic addition or change of a cryptographic definition to an LPAR (see Figure 10-7).

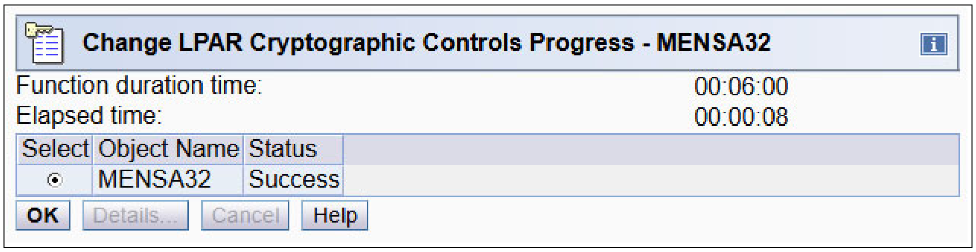

Figure 10-7 SE: Change LPAR Cryptographic Controls

Dynamic removal of the cryptographic definition

You can remove the cryptographic definition from an LPAR dynamically by using the Change LPAR Cryptographic Controls task. This section addresses the related issues and describes the procedure.

Complete the following steps:

1. Before you change the cryptographic settings by using the Change LPAR Cryptographic Controls task, verify that the cryptographic lists that you want to remove from an LPAR are offline (Standby). For more information about setting the cryptographic channel status, see "Configuring a Crypto Express8S online or offline on a logical partition" on page 212. If you try to remove the lists dynamically while they are online, the change fails and you receive the message that is shown in Figure 10-8.

   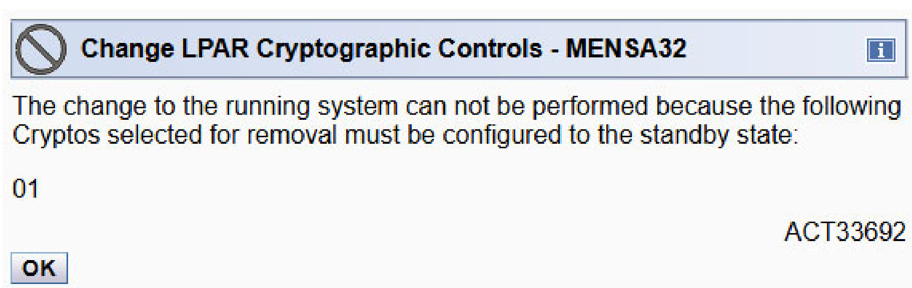

   Figure 10-8 SE: Change LPAR Cryptographic Controls - ACT33692

   Besides adding or changing cryptographic settings for an LPAR, you can remove the Control and Usage domains or Cryptographic Candidate lists for an LPAR from the Change LPAR Cryptographic Controls window (see Figure 10-6 on page 202).

   After clearing the definitions for an LPAR, remove a definition dynamically by clicking **Change Running System**. To save the new configuration to the image profile without changing the running system, select **Save to Profiles**. With **Save and Change**, the removal becomes concurrently active, and the removed cryptographic coprocessor is also not used for the next image activation.

2. When you remove the only definition of the cryptographic lists, the zeroize window opens (see Figure 10-9).

> **Consideration:** Because you cannot see all cryptographic information, including the usage domains for other LPARs, check the information in the **View LPAR Cryptographic Controls** window before you continue. For more information about zeroize, see “Reconfiguring the PCIe adapter type” on page 205.

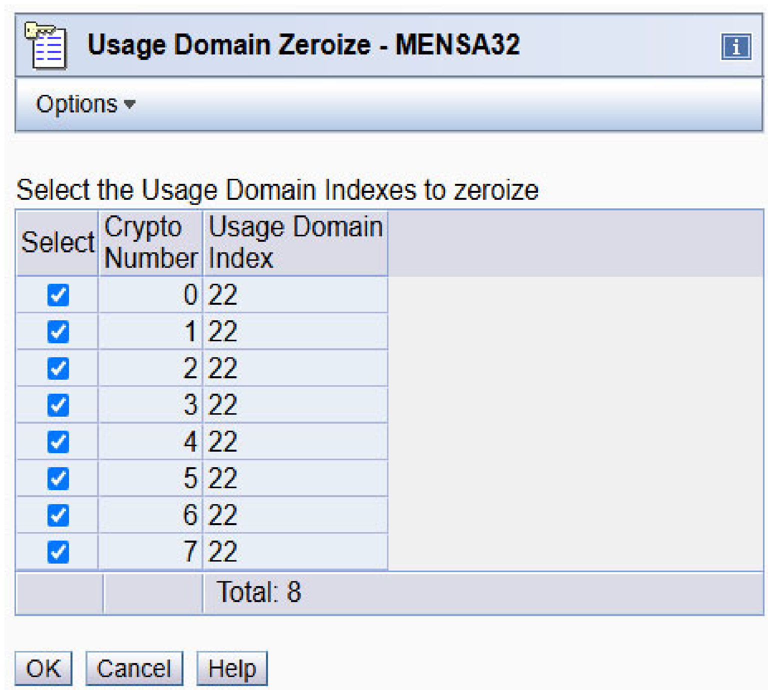

Figure 10-9 SE: Change LPAR Cryptographic Controls - Zeroize

3. In the confirmation window (see Figure 10-10 on page 205), click **OK** to dynamically change the cryptographic settings. You must also enter your user password to confirm this change. After processing, a status window indicates the result of the dynamic change of cryptographic definition to an LPAR.

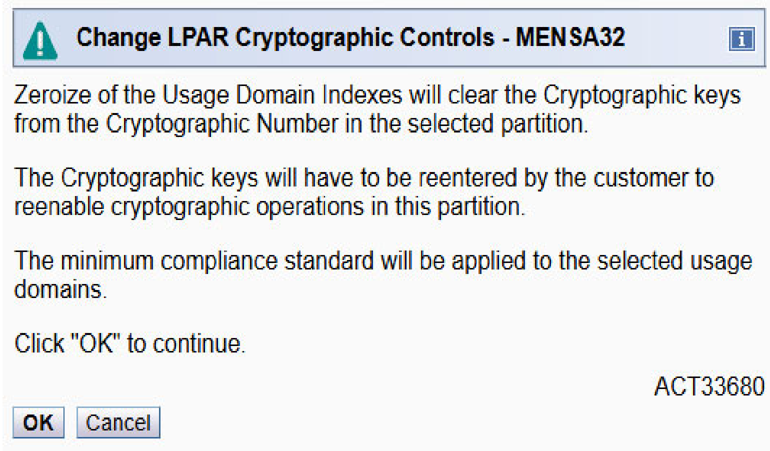

Figure 10-10 SE: Change LPAR Cryptographic Controls - ACT33680

Reconfiguring the PCIe adapter type

Each PCIe Crypto Express8S feature can be configured either as a coprocessor or as an accelerator. Each Crypto Express8S feature can be set in these configurations:

- CCA Coprocessor (CEX8C)
- PKCS #11 (EP11) Coprocessor (CEX8P)
- Accelerator (CEX8A)

Whether it is configured as a coprocessor or an accelerator, each PCIe Cryptographic adapter can be shared among 85 LPARs.

To reconfigure the cryptographic coprocessor, it must be offline to all LPARs. To put a cryptographic coprocessor online or offline on an LPAR requires using the Configure On/Off task, which is described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212. In the following example, we reconfigure the cryptographic coprocessor number 05, which is set offline to all LPARs.

Configuring a CCA coprocessor as an accelerator

During the installation of a Crypto Express8S feature, the PCIe Cryptographic adapter is configured by default as a CCA coprocessor. The reconfiguration is fully supported in LIC.

When a PCIe adapter is configured as a CCA coprocessor, it can still run accelerator functions, although slower than when configured as an accelerator. When it is configured as an accelerator, it cannot run coprocessor functions.

When a PCIe adapter is configured as an EP11 coprocessor, a TKE workstation is required for the management of the Crypto Express8S. For more information about configuring an EP11 coprocessor, see “Configuring a CCA coprocessor as an EP11 coprocessor” on page 209.

To reconfigure the PCIe adapter from coprocessor to accelerator, complete the following steps:

1. Select the CPC that has cryptographic coprocessor adapters that you want to reconfigure, and then click the **Cryptographic Configuration** task in the Configuration Group.
2. The reconfiguration is enabled only for PCIe adapters that are off. Therefore, be sure that the PCIe Cryptographic adapter status for that cryptographic coprocessor channel is unconfigured. If necessary, set the PCIe Cryptographic adapter to **Off** for all partitions that have it in their candidate list. To set the PCIe Cryptographic adapter to **Off**, use the

procedure that is described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212.

3. Select the number of the cryptographic coprocessor channel (see Figure 10-11) and click **Crypto Type Configuration**.

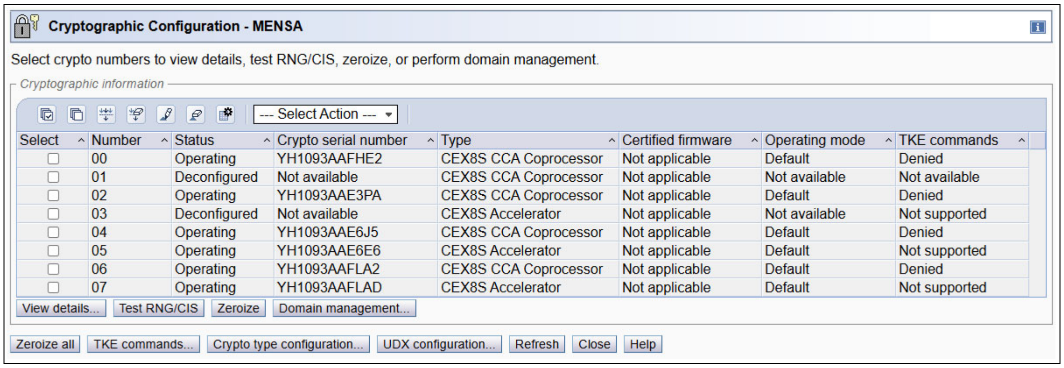

Figure 10-11 Cryptographic Configuration task (unconfigured)

4. Change the configuration for the cryptographic coprocessor adapter. The selected cryptographic coprocessor channel is configured as a coprocessor (see Figure 10-12). Select **Accelerator**.

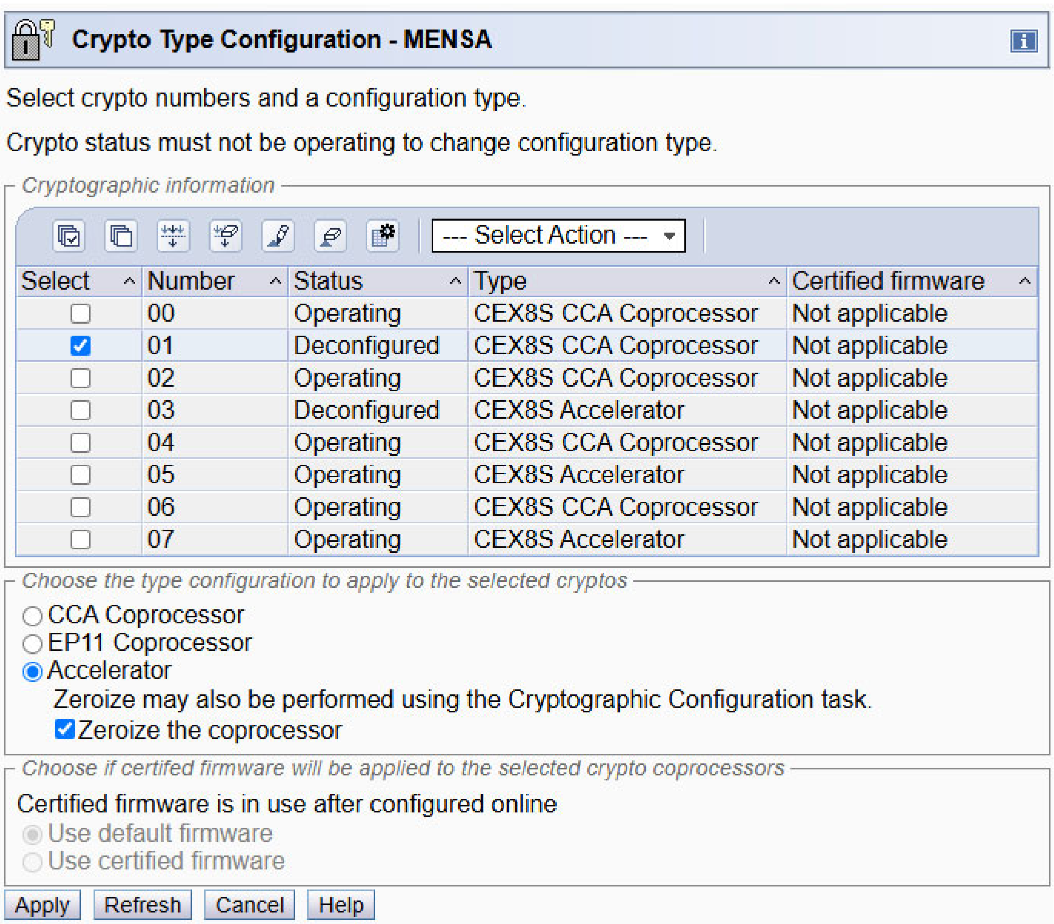

Figure 10-12 Crypto Type Configuration (CCA coprocessor to accelerator)

By selecting **Accelerator**, you can zeroize the selected coprocessor by also selecting **Zeroize the Coprocessor** on the Crypto Type Configuration window. Clear the **Zeroize the Coprocessor** checkbox and click **Apply**.

> **Important:** Zeroizing one or all cryptographic coprocessors clear their configuration data and all cryptographic keys. Zeroizing also erases configuration data from the SE hard disk drive (HDD) (for example, UDX files). Zeroize cryptographic coprocessors manually only when necessary, typically when the cryptographic coprocessor configuration data must be erased completely. In normal cases, be sure to clear the checkbox for each cryptographic channel.

5. Click **Yes** (see Figure 10-13).

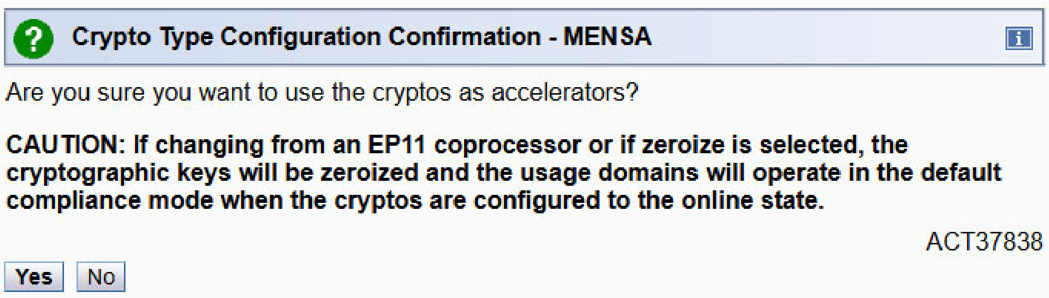

Figure 10-13 Crypto Type Configuration Confirmation for accelerator

6. Verify that your request was completed successfully. Click **OK**.
7. You are returned to the Crypto Type Configuration window. Click **Cancel**. You are returned to the Cryptographic Configuration window. Confirm that the target cryptographic channel changed to the cryptographic accelerator type. The Crypto Serial Number, Operating Mode, and TKE Commands should be “Not available” until the cryptography is set to Online again, as described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212.

   After you perform this task and go back to the Cryptographic Configuration window, where the information in Figure 10-14 appears.

> **Note:** UDX support is *not* available for a Crypto Express8S that is defined as an EP11 coprocessor and accelerator.

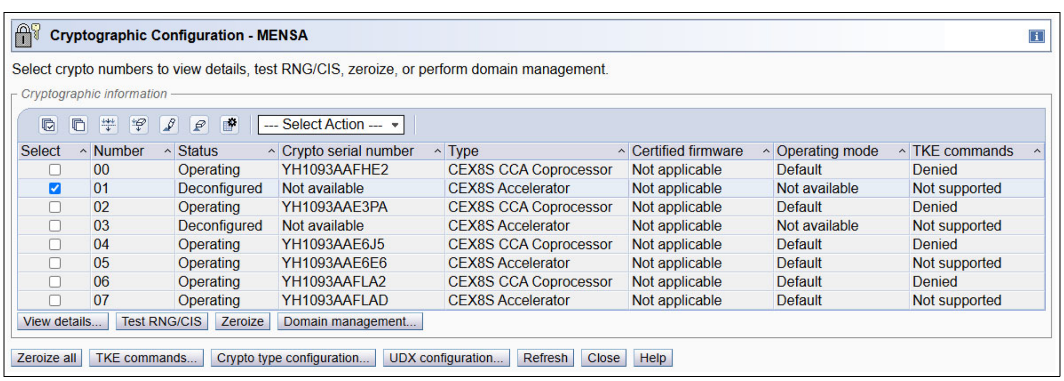

Figure 10-14 Cryptographic Configuration (Accelerator)

8. For more information, click **View Details** (see Figure 10-15).

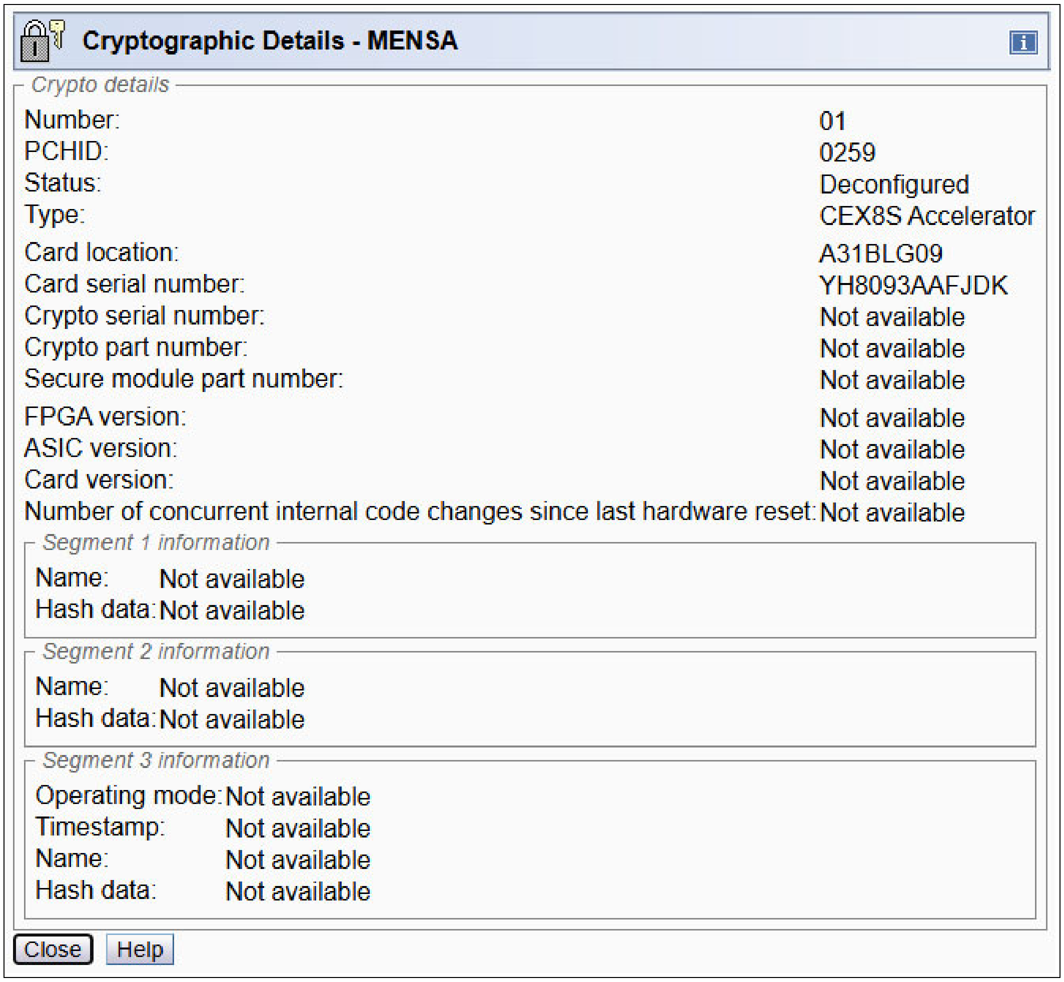

Figure 10-15 Cryptographic Details (Accelerator)

The Cryptographic Type is now a Crypto Express8S Accelerator. The adapter was not zeroized during the type-changing procedure.

All the information displayed as “Not available” will only be available after the crypto adapter is set online. See, “Setting a Crypto Express8S to an online state” on page 213.

The procedure for changing the type of the cryptographic configuration from a coprocessor to an accelerator is complete. To change the accelerator back to a coprocessor, the same procedure can be used, but select **Coprocessor** instead of **Accelerator**, as shown in Figure 10-12 on page 206.

The result of this change is shown in Figure 10-16 on page 209.

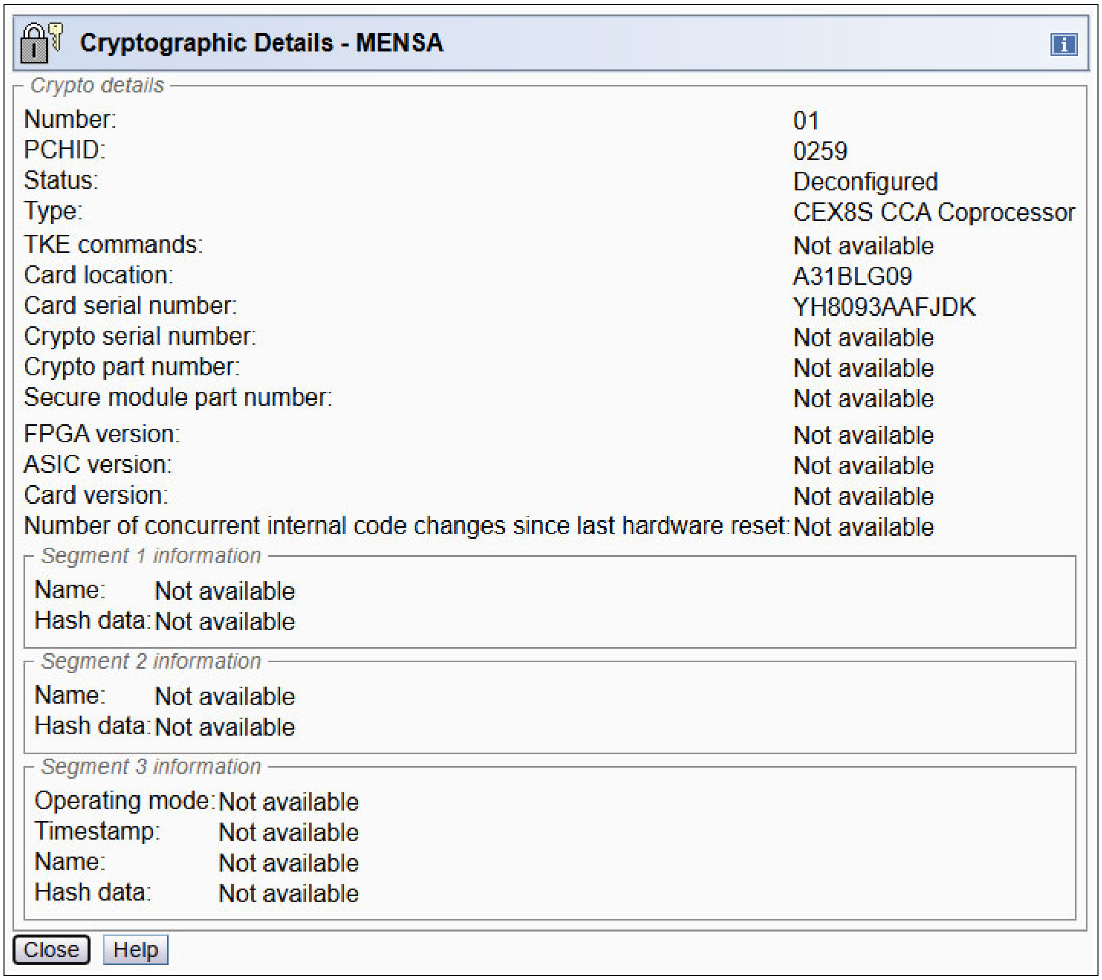

Figure 10-16 Cryptographic Details (CCA Coprocessor)

All the information displayed as “Not available” will only be available after the crypto adapter is set online. See, “Setting a Crypto Express8S to an online state” on page 213.

Configuring a CCA coprocessor as an EP11 coprocessor

To configure a CCA coprocessor as an EP11 coprocessor, complete the following steps:

1. Select the CPC that has cryptographic coprocessor adapters that you want to reconfigure and click **Cryptographic Configuration** in the CPC Configuration Group.
2. The reconfiguration is enabled only for PCIe adapters that are set to **Off**, so be sure the PCIe Cryptographic adapter status for that cryptographic coprocessor channel is unconfigured (Figure 10-11 on page 206). If necessary, set the PCIe Cryptographic adapter to **Off** for all partitions that have it in their candidate list. To set the PCIe Cryptographic adapter to **Off**, use the procedure that is described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212.
3. Select the number of the cryptographic coprocessor channel and click **Crypto Type Configuration**.

4. Change the configuration for the cryptographic coprocessor adapter. Select **EP11 Coprocessor** (see Figure 10-17), which by default automatically selects the **Zeroize the coprocessor** option. Click **Apply**.

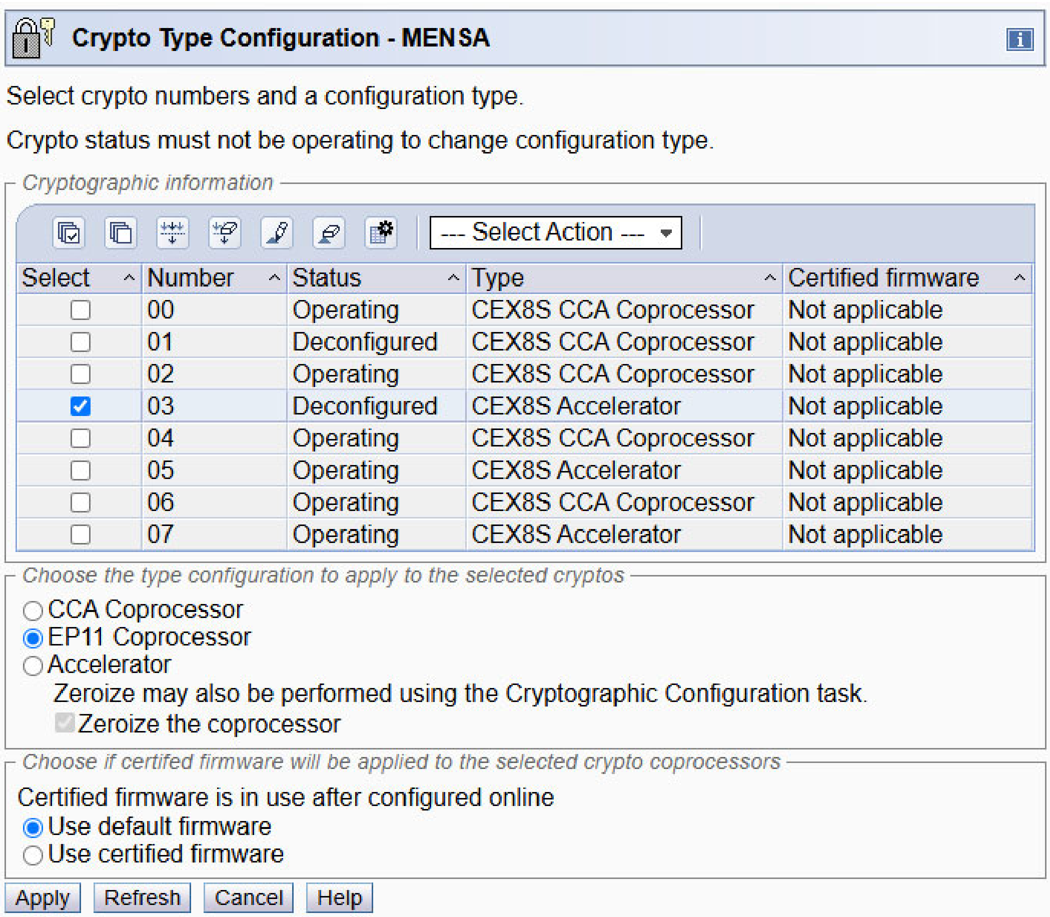

Figure 10-17 Crypto Type Configuration (CCA Coprocessor to EP11 Coprocessor)

5. Confirm your selection by clicking **Yes** (see Figure 10-18).

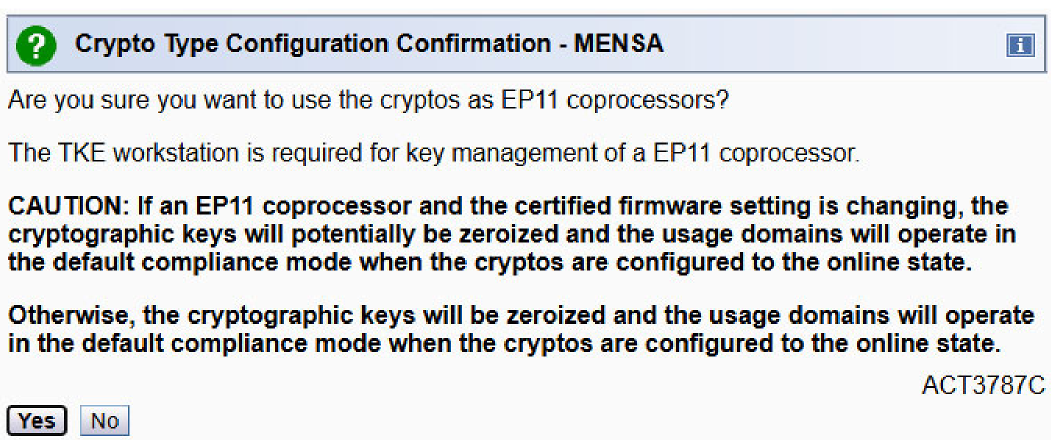

Figure 10-18 Crypto Type Configuration Confirmation for EP11 Coprocessor

6. Check that your request was completed successfully. Click **OK** to return to the Crypto Type Configuration window.

7. Click **Cancel** on the Crypto Type Configuration window to return to the Cryptographic Configuration window. Confirm that the target cryptographic channel changed to the EP11 Coprocessor type in the Cryptographic Configuration task window. The Crypto Serial Number, Operating Mode, and TKE Commands should be “Not available” until the cryptography is set to Online again, as described in “Configuring a Crypto Express8S online or offline on a logical partition” on page 212.

   After you complete this task and return to the Cryptographic Configuration window, the information in Figure 10-19 appears.

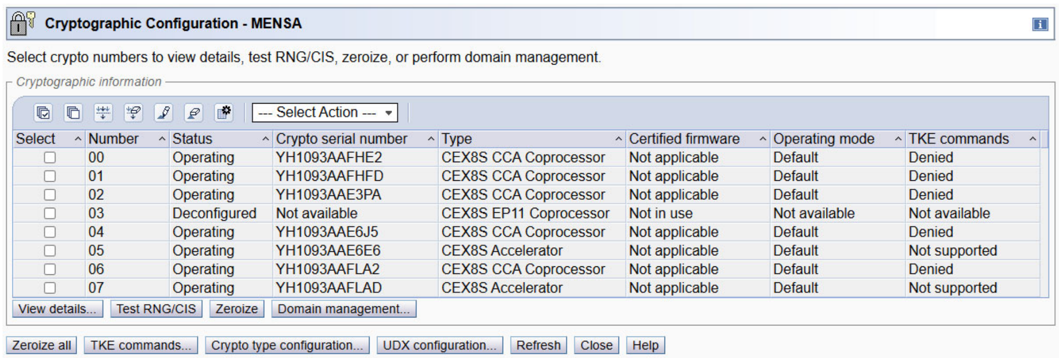

Figure 10-19 Cryptographic Configuration (EP11 Coprocessor)

Click **View Details** to display the detailed information, as shown in Figure 10-20 on page 212.

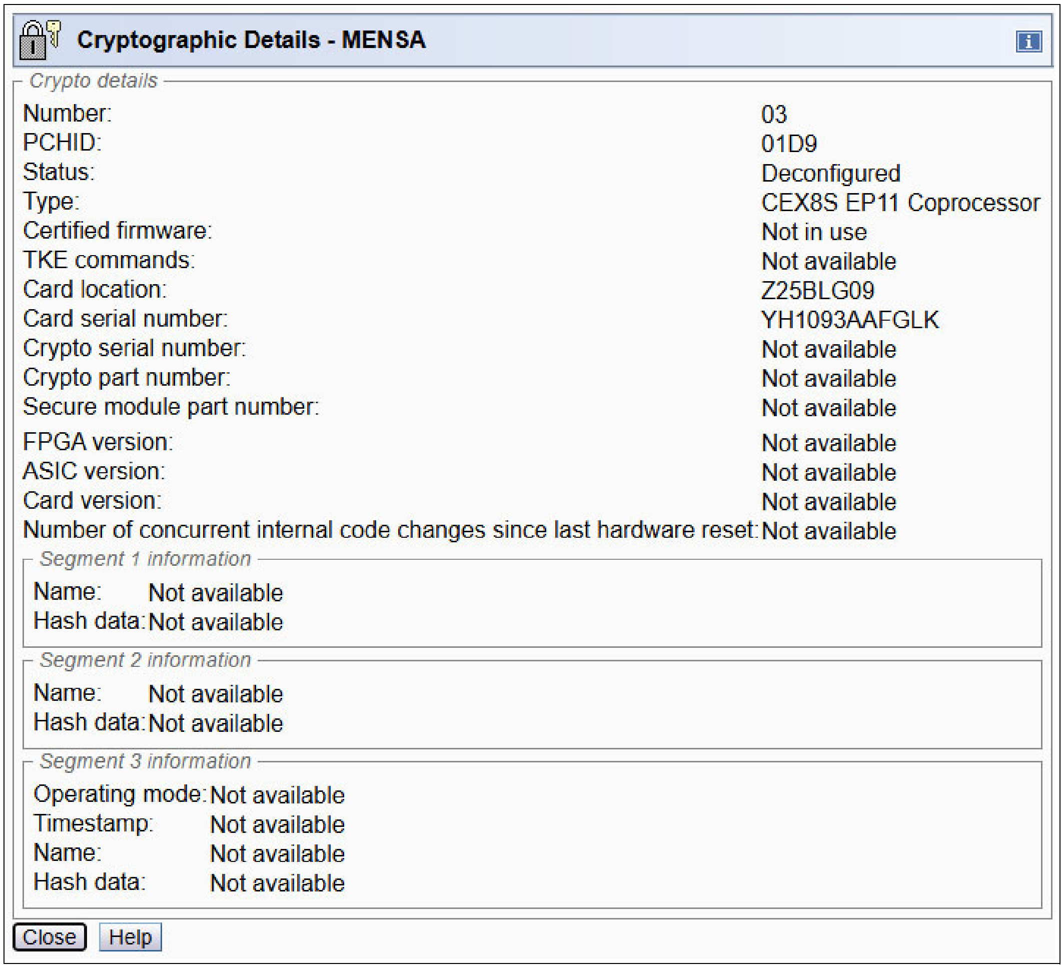

Figure 10-20 Cryptographic details (EP11 Coprocessor)

The Cryptographic Type is now a Crypto Express8S EP11 Coprocessor.

All the information displayed as “Not available” will only be available after the crypto adapter is set online. See, “Setting a Crypto Express8S to an online state” on page 213.

This step completes changing the type of the cryptographic configuration from a CCA Coprocessor to an EP11 coprocessor. To change the configuration back to a CCA Coprocessor, the same procedure can be used, but select **CCA Coprocessor** instead of **EP11 Coprocessor**.

You can also switch the configuration mode from Accelerator to EP11 Coprocessor and from EP11 to Accelerator by using the same process, but selecting **Accelerator** or **EP11 Coprocessor** as required.

Requirement: To manage a Crypto Express8S feature that is configured as an EP11 coprocessor, the TKE workstation is required.

Configuring a Crypto Express8S online or offline on a logical partition

For some changes to the cryptographic settings to the LPAR, you must configure the Crypto Express8S online or offline. If you can reactivate (DEACTIVATE and ACTIVATE) the image for the LPARs whose cryptographic online lists were updated, this dynamic operation is not needed.

Setting a Crypto Express8S to an online state

To set a Crypto Express8S online, complete the following steps:

1. From the SE, select the **System Management** function.
2. Select the server, click **Partitions**, and then select the LPAR.
3. Click the **Cryptos** selection for the target LPAR.
4. In the Cryptos page, select the Crypto IDs to be changed. Figure 10-21 shows that on server MENSA to LPAR MENSA32, eight cryptographic coprocessors are defined: Four CCA coprocessors (ID 00, 02, 04, and 06, PCHID 0258, 01D8, 0118, and 0220), and four accelerators (IDs 01, 03, 05, and 07, PCHIDs 0259, 01D9, 0119 and 0221). The IDs 00, 01, 02, 03, 04, 06, and 07 are online, and the Accelerator ID 05 is offline. We also want to set up this accelerator online.

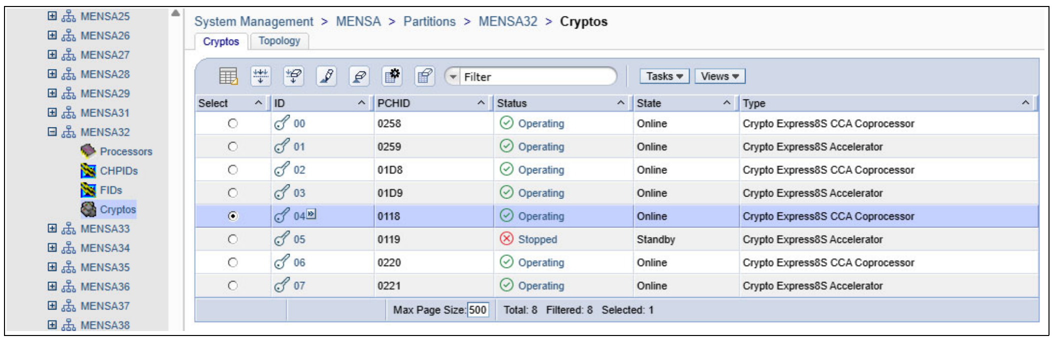

Figure 10-21 System Management: LPAR Crypto Selection Standby

5. Select the cryptographic coprocessor, and then select **Tasks** → **Crypto Service Operations** → **Configure On/Off task** (see Figure 10-22 on page 213). This task controls the online or offline (standby) state of a cryptographic processor for LPARs that are defined in the cryptographic processor's candidate list.

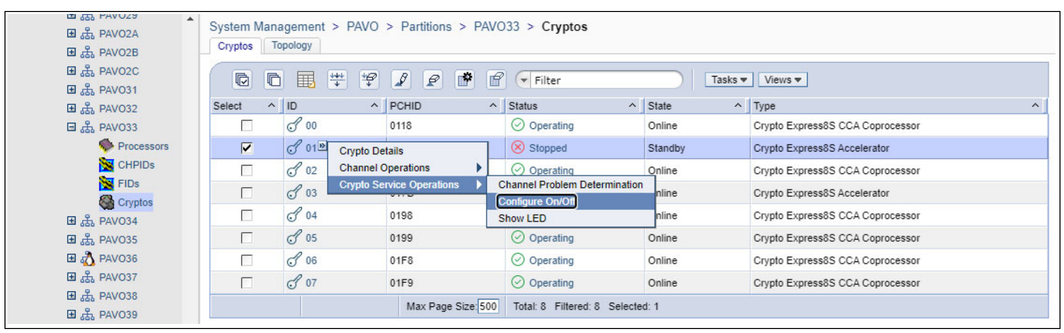

Figure 10-22 Crypto Service Operations: LPAR Crypto Selection Configure Online

6. Select the cryptographic coprocessor channel number that you want, and then select **Select Action** → **Toggle** to switch from Standby to Online (see Figure 10-23). If you want multiple cryptographic channels concurrently, select **Toggle All On**.

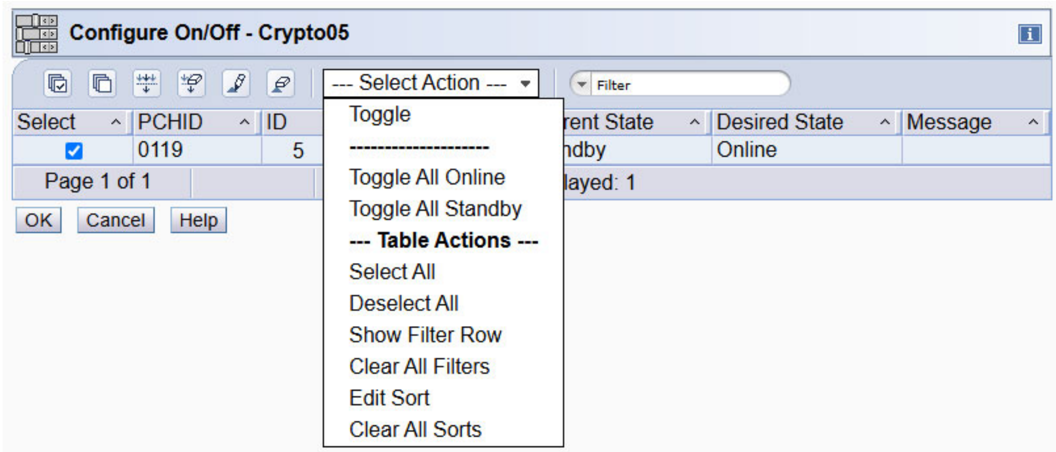

Figure 10-23 Configure On/Off (Standby)

7. After confirming that your requested cryptographic coprocessor channel is set to the state of Online, click **OK** (see Figure 10-24).

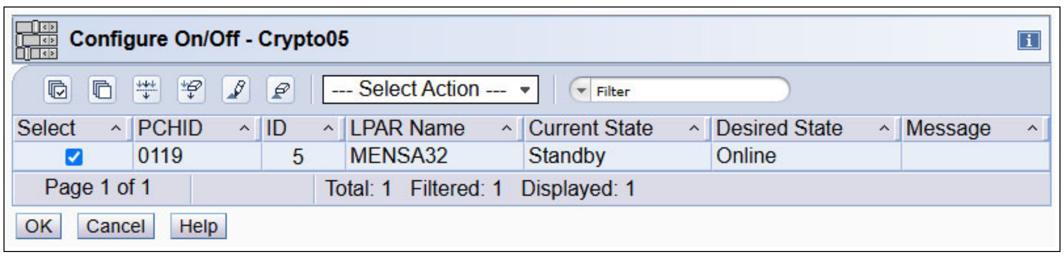

Figure 10-24 Configure On/Off (Standby to Online)

8. Confirm that your request is completed (see Figure 10-25). Click **OK**.

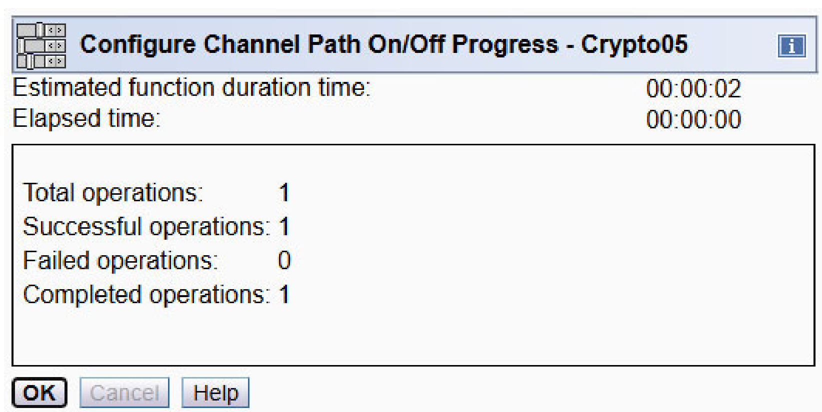

Figure 10-25 Configure On/Off (Standby to Online) completed

9. You can now view the contents of the Cryptos window of the LPAR to confirm that the cryptographic channels are now in the Operating status (see Figure 10-26).

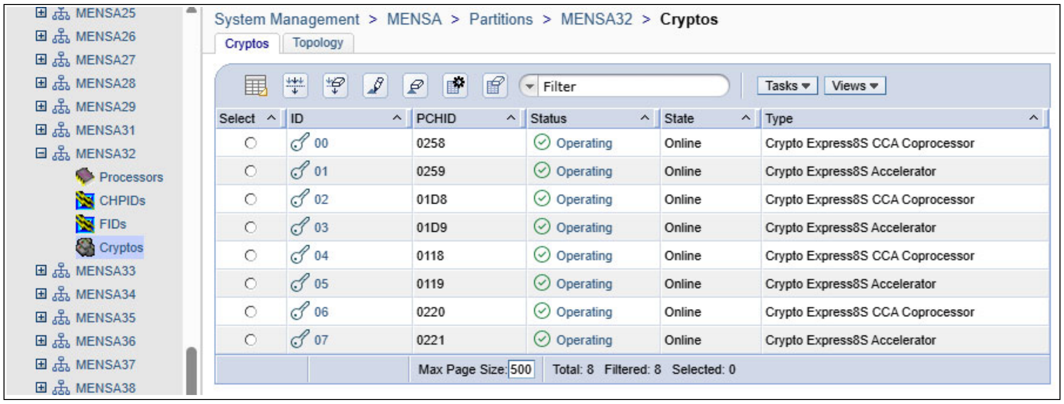

Figure 10-26 System Management: LPAR Crypto Selection Online

Changing a cryptographic channel to standby (offline) status

To change the cryptographic channel status, complete the following steps:

1. Select the LPAR whose Crypto IDs you want to change to Standby. For example, select the accelerator (03) that is in an online state. Select **Tasks** → **Crypto Service Operations** → **Configure On/Off task** (see Figure 10-27).

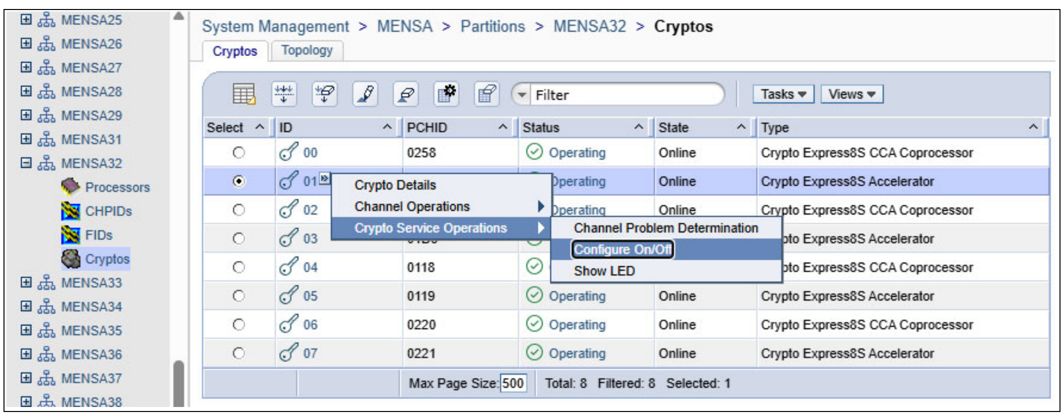

Figure 10-27 Crypto Service Operations: LPAR Crypto Selection, Configure Offline

2. Select the cryptographic coprocessor channel number that you want, and select **Select Action** → **Toggle All Standby** to switch from Online to Standby (see Figure 10-28 on page 216).

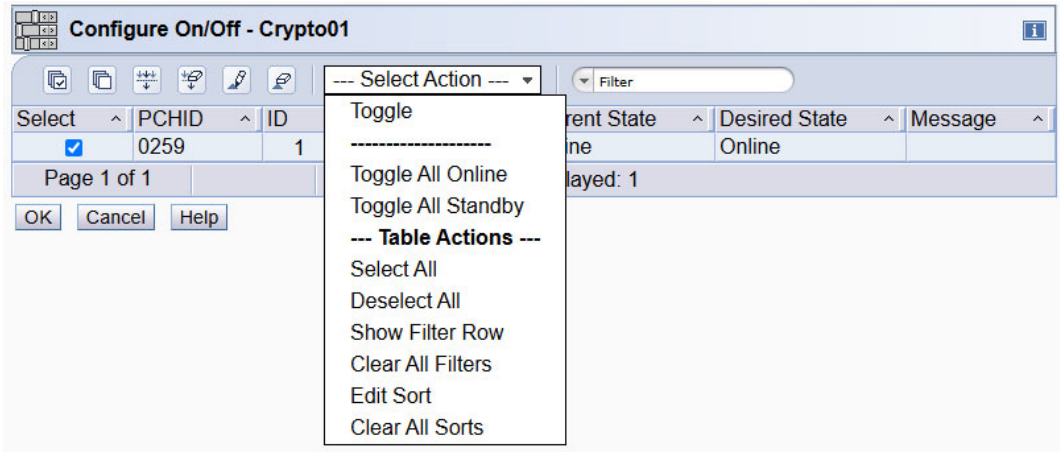

Figure 10-28 Configure On/Off (Online)

3. After you confirm that the state for your requested cryptographic channel is Standby, click **OK** (see Figure 10-29).

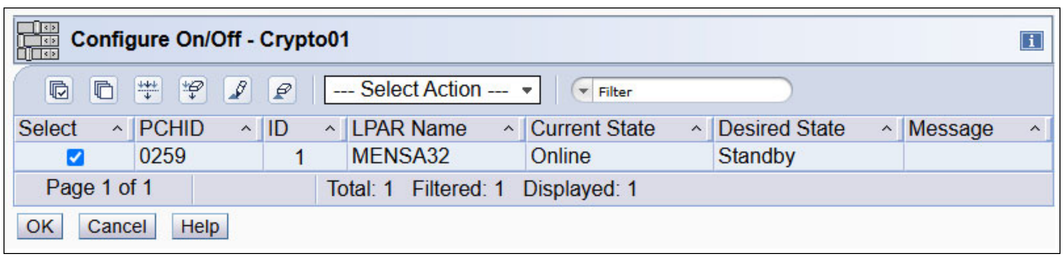

Figure 10-29 Configure On/Off (Online to Standby)

4. Because taking a cryptographic coprocessor offline can be disruptive to your application, a confirmation is required. The task must be confirmed by entering the user password (see Figure 10-30).

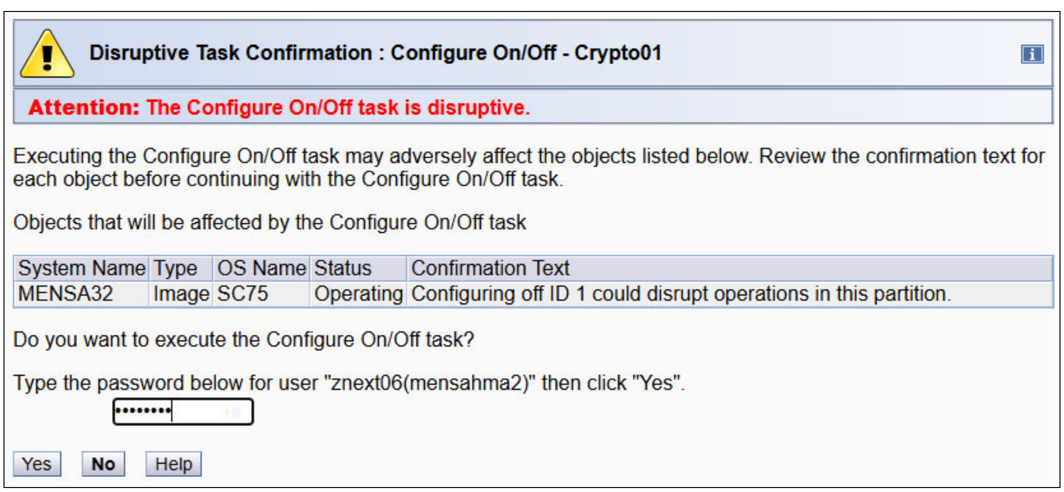

Figure 10-30 Configure On/Off (Online to Standby): Confirmation

5. Confirm that your request is completed (see Figure 10-31). Click **OK**.

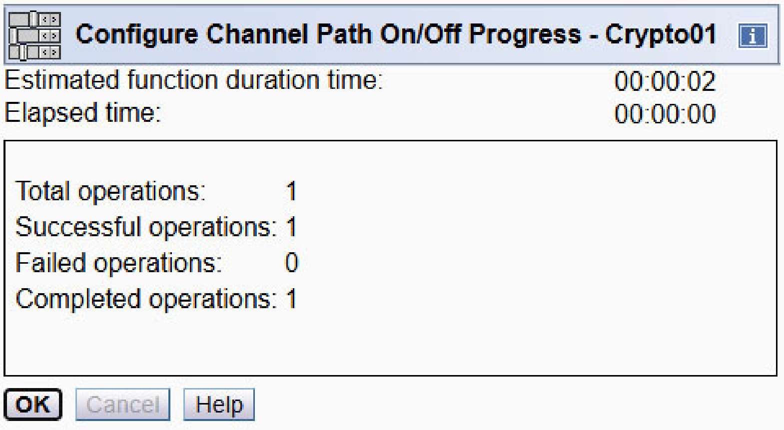

Figure 10-31 Config On/Off (Online to Standby) completed

10.1.4 Handling cryptographic coprocessors by using ICSF

ICSF provides an Interactive System Productivity Facility (ISPF) Coprocessor Management panel where you can display or change the status (Active or Deactivate) of cryptographic coprocessors. This action affects only the coprocessor status of ICSF, and has no effect on the Online/Standby hardware status that is displayed on the IBM z17 SE.

It is not the purpose of this section to show how to create, load, and manage keys in the cryptographic adapters. For more information, see the ICSF literature. This section shows only how to activate and deactivate a cryptographic coprocessor and display the hardware status.

Complete the following steps:

1. From the ICSF main panel (see Figure 10-32 on page 218), select option 1 to open the ICSF Coprocessor Management panel.

```
HCR77E0 -------------- Integrated Cryptographic Service Facility -------------
OPTION ===> 1
System Name:  SC75                           Crypto Domain: 22
Enter the number of the wanted option.

  1  COPROCESSOR MGMT -  Management of Cryptographic Coprocessors
  2  KDS MANAGEMENT   -  Master key set or change, KDS Processing
  3  OPSTAT           -  Installation options
  4  ADMINCNTL        -  Administrative Control Functions
  5  UTILITY          -  ICSF Utilities
  6  PPINIT           -  Pass Phrase Master Key/KDS Initialization
  7  TKE              -  TKE PKA Direct Key Load
  8  KGUP             -  Key Generator Utility processes
  9  UDX MGMT         -  Management of User-Defined Extensions

     Licensed Materials - Property of IBM
     5655-ZOS Copyright IBM Corp. 1989, 2023.
     US Government Users Restricted Rights - Use, duplication or
     disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Press ENTER to go to the selected option.
Press END   to exit to the previous menu.
```

Figure 10-32 Integrated Cryptographic Support Facility main panel

Cryptographic coprocessors that are configured on the partition are listed in the ICSF Coprocessor Management panel (see Figure 10-33).

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 8 of 8
COMMAND ===>                                                SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R, S, and V. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS                AES   DES   ECC   RSA   P11
  -------   --------  --------------------  ---   ---   ---   ---   ---
_   8C00    93AAFHE2  Active                 A     A     A     A
_   8A01    N/A       Active
_   8C02    93AAE3PA  Active                 A     A     A     A
_   8A03    N/A       Active
_   8C04    93AAE6J5  Active                 A     A     A     A
_   8A05    N/A       Active
_   8C06    93AAFLA2  Active                 A     A     A     A
_   8A07    N/A       Active
****************************** Bottom of data ********************************
```

Figure 10-33 ICSF Coprocessor Management

When a coprocessor is configured offline to the LPAR from the SE (Standby status), it is shown as `Offline` in the ICSF Coprocessor Management panel (see Figure 10-34).

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 8 of 8
COMMAND ===>                                                SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R, S, and V. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS               AES   DES   ECC   RSA   P11
  -------   --------  -------------------- ---   ---   ---   ---   ---
_   8C00    93AAFHE2  Active                A     A     A     A
_   8A01    N/A       Active
_   8C02    93AAE3PA  Active                A     A     A     A
_   8A03    N/A       Offline
_   8C04    93AAE6J5  Active                A     A     A     A
_   8A05    N/A       Active
_   8C06    93AAFLA2  Active                A     A     A     A
_   8A07    N/A       Active
******************************* Bottom of data ********************************
```

Figure 10-34 ICSF Coprocessor Management (Candidate only - Standby)

A cryptographic coprocessor becomes visible to ICSF Coprocessor Management when the coprocessor number is part of the partition candidate list and the coprocessor is first brought online to the partition in either of these ways:

– At the time the partition is activated, if the coprocessor is installed and the coprocessor number is part of the partition Online list.

– When the coprocessor is first configured online to the partition by using the Config On/Off task from the SE Workplace.

2. In the list (see Figure 10-35), enter `A` or `D` to switch a coprocessor status to `Active` or `Deactivated`.

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 8 of 8
COMMAND ===>                                                SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R, S and V. See the help panel for details.

CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS               AES   DES   ECC   RSA   P11
  -------   --------  -------------------- ---   ---   ---   ---   ---
_   8C00    93AAFHE2  Active                A     A     A     A
_   8A01    N/A       Active
_   8C02    93AAE3PA  Active                A     A     A     A
d   8A03    N/A       Active
_   8C04    93AAE6J5  Active                A     A     A     A
_   8A05    N/A       Active
_   8C06    93AAFLA2  Active                A     A     A     A
_   8A07    N/A       Active
******************************* Bottom of data ********************************
```

Figure 10-35 ICSF Coprocessor Management (Online)

When a coprocessor is deactivated through ICSF (see Figure 10-36), it cannot be used by applications that run in that system image. The EP11 coprocessor configuration requires a TKE workstation.

Generally, deactivate an active coprocessor from the ICSF Coprocessor Management panel before it is configured off from the SE.

Note: If you do not deactivate the coprocessor first, some jobs might not be rerouted correctly.

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 8 of 8
COMMAND ===>                                                  SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R, S and V. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS                AES   DES   ECC   RSA   P11
  -------   --------  --------------------  ---   ---   ---   ---   ---
_   8C00    93AAFHE2  Active                 A     A     A     A
_   8A01    N/A       Active
_   8C02    93AAE3PA  Active                 A     A     A     A
_   8A03    N/A       Deactivated
_   8C04    93AAE6J5  Active                 A     A     A     A
_   8A05    N/A       Active
_   8C06    93AAFLA2  Active                 A     A     A     A
_   8A07    N/A       Active
******************************* Bottom of data ********************************
```

Figure 10-36 SF Coprocessor Management (Deactivated)

The Active/Deactivated status that is viewed from ICSF Coprocessor Management does not change the Online/Standby status that is set from the IBM z17 SE.

Figure 10-37 shows ICSF Coprocessor Hardware Status panel for adapters 8C00 and 8C02.

```
---------------------- ICSF - Coprocessor Hardware Status ---------------------
COMMAND ===>                                                  SCROLL ===>
                                                             CRYPTO DOMAIN: 22

REGISTER STATUS                  COPROCESSOR 8C00       COPROCESSOR 8C02
                                                                  More:     +
 Crypto Serial Number          : 93AAFHE2               93AAE3PA
 Status                        : ACTIVE                 ACTIVE
 PCI-HSM Compliance Mode       : INACTIVE               INACTIVE
 Compliance Migration Mode     : INACTIVE               INACTIVE
 AES Master Key
   New Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Current Master Key register : VALID                  VALID
     Verification pattern      : 1300CF50ECF4DEA6       1300CF50ECF4DEA6
 DES Master Key
   New Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Current Master Key register : VALID                  VALID
     Verification pattern      : 29069A18A233405A       29069A18A233405A
ECC Master Key
   New Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
   Current Master Key register : VALID                  VALID
     Verification pattern      : AFDEEE7A0AB38568       AFDEEE7A0AB38568
 RSA Master Key
   New Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
                               :
   Old Master Key register     : EMPTY                  EMPTY
     Verification pattern      :
                                :
   Current Master Key register : VALID                  VALID

Press ENTER to refresh the hardware status display.
```

Figure 10-37 ICSF Coprocessor Hardware Status

Help information from the ICSF Coprocessor Management panel (see Figure 10-38 and Figure 10-39 on page 223) describes valid actions and status information for each type of cryptographic coprocessor.

```
----------------------- Help for Coprocessor Management -----------------------
COMMAND ===>
Press enter to page through this help.
                                                                   More:     +
The Coprocessor Management panel displays the status of all cryptographic
coprocessors installed.  Select the coprocessors to be processed.

Prefix     Type of cryptographic coprocessor        Valid action characters
------     ---------------------------------        -----------------------
  5A        Crypto Express5 Accelerator              a, d
  5C        Crypto Express5 CCA coprocessor          a, d, e, k, r, s, v
  5P        Crypto Express5 PKCS #11 coprocessor     a, d, r, s, v
  6A        Crypto Express6 Accelerator              a, d
  6C        Crypto Express6 CCA coprocessor          a, d, e, k, r, s, v
  6P        Crypto Express6 PKCS #11 coprocessor     a, d, r, s, v
  7A        Crypto Express7 Accelerator              a, d
  7C        Crypto Express7 CCA coprocessor          a, d, e, k, r, s, v
  7P        Crypto Express7 PKCS #11 coprocessor     a, d, r, s, v
  8A        Crypto Express8 Accelerator              a, d
  8C        Crypto Express8 CCA coprocessor          a, d, e, k, r, s, v
  8P        Crypto Express8 PKCS #11 coprocessor     a, d, r, s, v

Action characters:  (entered on the left of the coprocessor number)
 'a'      Makes available a coprocessor previously deactivated by a 'd'.
 'd'      Makes a coprocessor unavailable.
 'e'      Selects the coprocessor for master key entry.
 'k'      Selects the coprocessor for operational key load.
 'r'      Causes the coprocessor domain role to be displayed.
 's'      Causes complete hardware status to be displayed.
 'v'      Causes the coprocessor default role to be displayed with offsets.
The action character 'e' cannot be combined with any other action characters.
The action character 'k' may be specified on only one coprocessor.
The action characters' may not be specified for both CCA and
  PKCS #11 coprocessors at the same time.

Status:
  - Active:               The feature is available for work.
  - Offline:              The feature is installed but not available to ICSF.
  - Deactivated:          The feature has been deactivated (see action
                            characters)
  - Busy:                 The feature is temporarily busy.
  - Hardware error:       The feature has been stopped.
  - Disabled by TKE:      The feature has been removed from service by a TKE
                            workstation.
  - Master key incorrect: At least one master key is incorrect.
  - Being reconfigured:   An error has been detected and being checked by the
                            configuration task
```

Figure 10-38 Help for Coprocessor Management (part 1 of 2)

```
  - Initializing stage 1: The feature has been detected by the configuration
                            task. No status is available.
  - Initializing stage 2: The feature is being reset by the configuration
                            task. No status is available.
  - Initializing stage 3: The feature is being readied by configuration
                            task. No status is available.
  - Unknown response:     The feature has returned a return/reason code
                            combination that ICSF does not recognize.
  - Hung User on Feature: The feature is not responding. One or more users
                            hold the feature latch. If this problem persists
                            please dump and contact IBM service.
                            You will need to recycle ICSF to reclaim use of
                            the feature.
  - Bad feature response: An unexpected response was received from a feature.
  - Retry limit reached:  Initialization of the feature failed.
  - Unknown feature type: A feature has a type that is not recognized by ICSF.
  - Repeat failures:      A feature has experienced repeated failures
                            and recovered.  The feature is made inactive and
                            will require manual intervention to cause ICSF
                            to attempt to use it again.

Cryptographic Coprocessor Master Key State:
    A:  Master key Verification Pattern matches the keystore (CKDS, PKDS, or
        TKDS) and the master key is available for use
    C:  Master key Verification Pattern matches the keystore, but the master
        key is not available for use
    E:  Master key Verification Pattern mismatch for keystore or for P11, no
        TKDS was specified in the options data set
    I:  The Master key Verification Pattern in the keystore is not set,
        so the contents of the Master key are Ignored
    U:  Master key is not initialized
    -:  Not supported
     :  Not applicable

    F3 = END HELP
```

Figure 10-39 Help for Coprocessor Management (part 2 of 2)

10.2 Virtual Flash Memory

This section provides information about the configuration of the Virtual Flash Memory (VFM) feature on the IBM z17 ME1. It covers the following topics:

- VFM overview
- Planning for VFM configuration
- Configuring VFM
- VFM management

10.2.1 VFM overview

VFM offers up to 6.0 TB on IBM z17 ME1, ordered in 512 GB increments.

VFM is designed to help improve the availability and handling of paging workload spikes when running on a supported z/OS version. With this support, z/OS helps to improve system availability and responsiveness by using VFM across transitional workload events, such as market openings and diagnostic data collection. z/OS is also designed to help improve processor performance by supporting middleware exploitation of pageable large (1 MB) pages.

Using VFM can help availability by reducing latency from paging delays that can occur at the start of the workday or during other transitional periods. It also helps to eliminate delays that can occur when collecting diagnostic data during failures.

With CFLEVEL 26, which comes with the IBM z17 Driver Level 61, the support for use of Virtual Flash Memory for Coupling Facility images has been removed. If a client was using Virtual Flash Memory as an "overflow" mechanism for MQ Shared Queue structures, the client should plan to use other alternatives such as using larger Coupling Facility structures, using IBM MQ shared message data sets (SMDS), or using MQ techniques to offload MQ message data to Db2.

VFM can help organizations meet their most demanding service-level agreements and compete more effectively. VFM is simple to configure to provide rapid time to value.

For more information about the VFM feature, see *IBM z17 (9175) Technical Guide*, SG24-8579.

10.2.2 Planning for VFM configuration

For planning considerations, see “Planning considerations for Virtual Flash Memory” on page 30.

10.2.3 Configuring VFM

The assignment of VFM to LPARs is exclusively done with the definitions in the image activation profiles.

> **Note:** Unlike the Flash Express cards, the allocation of VFM to LPARs cannot be altered on an activated LPAR. So, the **Manage Flash Allocation** selection on the HMC is not supported for IBM z15 T01, IBM z15 T02, IBM z16 A01, IBM z16 A02, IBM z16 AGZ and IBM z17 ME1 systems.

Consider the following items when you allocate VFM to a partition:

- When an allocation is first defined, you must set the initial and maximum allocation in 16 GB increments.
- A storage-class memory (SCM) allocation is put online to the z/OS image that is assigned to the partition at IPL time unless the z/OS image is configured not to do so.
- z/OS enables more memory to be configured online, up to the maximum memory that is defined in this window or up to the maximum VFM that is available and not used by other LPARs.

- Minimum amounts are allocated from the available pool, so they cannot be overallocated.
- Maximum amounts can be overallocated up to the VFM LICCC value of the IBM z17 ME1.
- Maximum amounts must be greater than or equal to the initial amounts.

To allocate VFM to a partition, select the LPAR on the HMC and select **Operational Customization** → **Customize/Delete Activation Profiles**. Then, select the image profile and click **Customize profile**. The Initial and the Maximum values for the VFM are specified on the **Storage** tab. This configuration is shown in Figure 10-40.

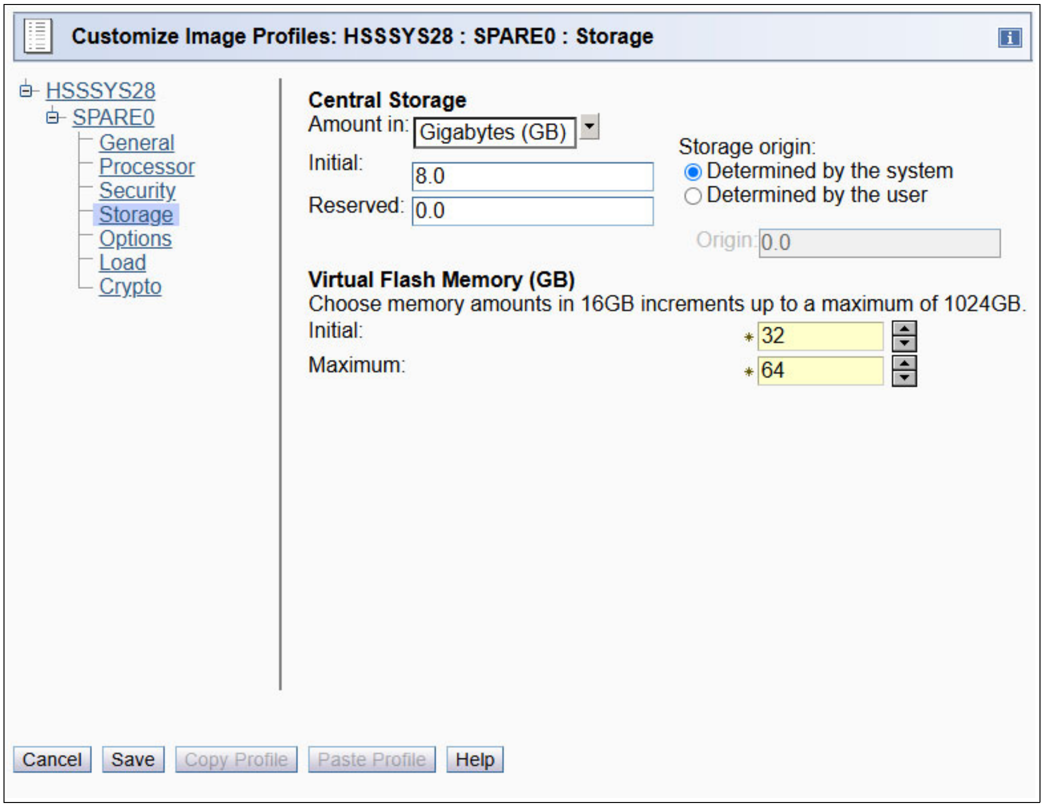

Figure 10-40 Virtual Flash Memory allocation in an image profile

In this example, the initial value is set to 32 GB, and the maximum value is set to 64 GB. The IBM z17 ME1 system has two VFM features that are installed, which enables a maximum of 1024 GB that is allocated to the LPAR.

These definitions do not change the settings of a running LPAR. They are used only for the activation of the LPAR. A newly activated LPAR starts with the specified amount of initial VFM.

If in the image activation profile of an LPAR an amount of initial VFM storage greater than the available amount of deallocated VFM on the IBM z17 ME1 is specified, the activation of this LPAR fails with message ACTZ01EB, as illustrated in Figure 10-41.

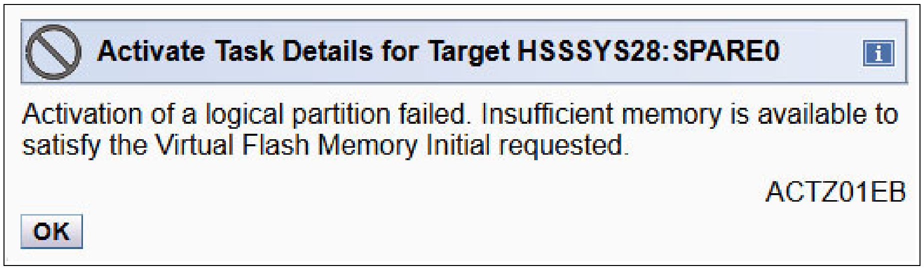

Figure 10-41 Insufficient VFM available: ACTZ01EB

10.2.4 VFM management

The memory allocation of an IBM z17 ME1 system is shown on the SE in the Storage Information window. To view it, click the server and select **Operational Customization** → **Storage Information**. The Base System Storage Allocation window opens, as shown in Figure 10-42.

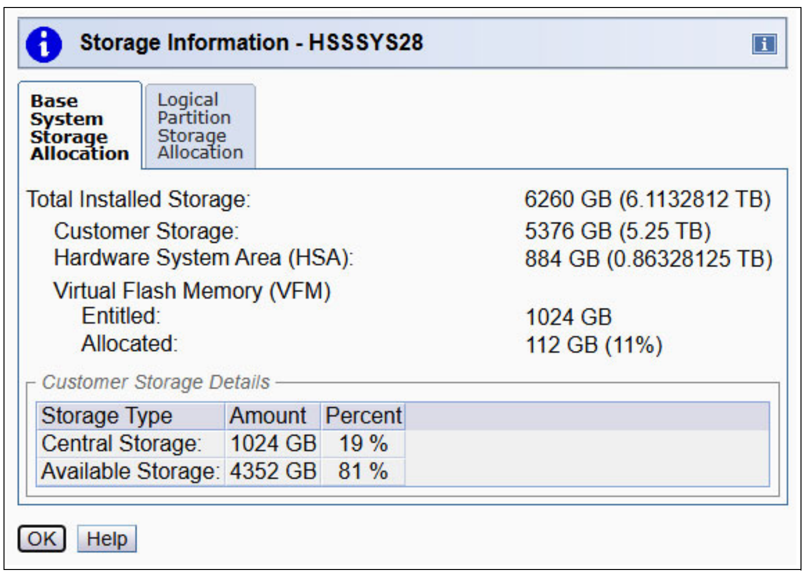

Figure 10-42 Storage Information: Base System Storage Allocation

The IBM z17 ME1 in our example (HSSSYS28) has 1024 GB installed, of which 112 GB is allocated to activated LPARs.

The Logical Partition Storage Allocation window shows the VFM allocation of the LPARs. For every LPAR, the initial and the maximum amount of VFM are listed (which were specified in the image activation profile), and the currently allocated amount, as shown in Figure 10-43.

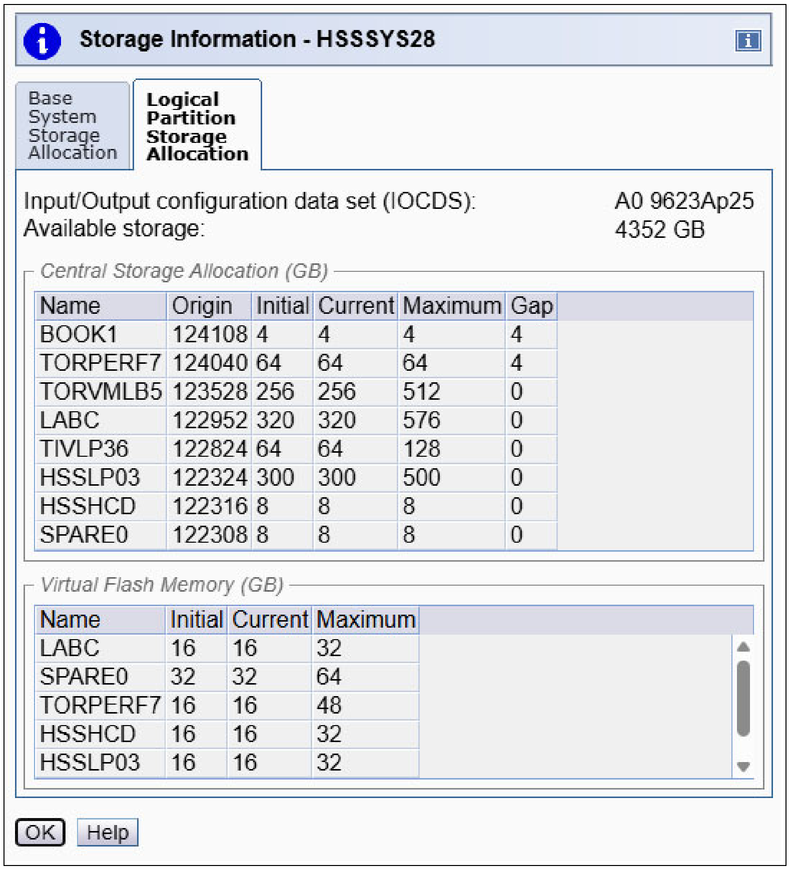

Figure 10-43 Storage Information: Logical Partition Storage Allocation

In z/OS for VFM, the **PAGESCM** parameter is supported in `IEASYSxx`. The syntax is shown in Example 10-1. This parameter determines whether and how much SCM is made available to an LPAR at IPL time.

Example 10-1 PAGESCM parameter

```
PAGESCM={xxxxxxM                          }
        {xxxxxxG                          }
        {xxT                              }
        {ALL                              }
        {NONE                             }
        {0                                }
```

This parameter specifies the minimum amount of SCM that should be made available for use as auxiliary storage. The system reserves this amount of SCM during IPL for subsequent use as auxiliary storage. Extra SCM is allocated on an as-needed basis if usage of this initial amount of SCM is exceeded.

You can specify the following value ranges for the **PAGESCM** parameter to reserve SCM for paging at IPL:

`xxxxxxM` Specifies the amount of SCM to reserve for paging at IPL in megabytes. This value can be 1 - 6 decimal digits.

xxxxxxG Specifies the amount of SCM to reserve for paging at IPL in gigabytes. This value can be 1 - 6 decimal digits.

xxT Specifies the amount of SCM to reserve for paging at IPL in terabytes. This value can be 1 - 2 decimal digits. The maximum amount of SCM supported for paging is 16 TB.

`ALL` Reserves all SCM for paging at IPL.

`NONE` SCM is not used for paging. This parameter remains in effect until the next IPL.

0 | 0M | 0G | 0T Indicates that no SCMs are reserved for paging at IPL. Instead, SCM is allocated as needed based on paging demand.

`Default value` ALL.

`Associated parmlib member`:None.

The **`CONFIG SCM`** command is used to set SCM online or offline to an LPAR (Example 10-2).

Example 10-2 CONFIG SCM

```
CONFIG SCM(ddddddddM|G|T),ONLINE|ON
CONFIG SCM(ddddddddM|G|T),OFFLINE|OFF
CONFIG SCM(scm_ranges),OFFLINE|OFF
```

The system reconfigures SCM both logically and physically. To bring SCM online, a number must be specified. To take SCM offline, a range of starting and ending addresses of the SCM blocks must be specified.

The command has the following values:

ddddddddM|G|T The amount of SCM to be reconfigured. Specify up to eight decimal digits followed by a multiplier (M = megabytes, G = gigabytes, or T = terabytes) for this amount. Check your processor configuration for the supported SCM increment sizes. The value for *dddddddd* must be a multiple of the SCM increment size (usually 2, 4, or 8), and cannot exceed `16T`.

Instead of specifying a decimal amount, you can alternatively specify a hexadecimal amount, with or without a multiplier, in the following format:

`X'xxxxxx'`

For example:

`X'123456789A00000'`

`X'123'M`

You can use underscore characters in any hexadecimal specification for more clarity. Underscore characters in the specification are ignored during processing.

> **Attention:** If you take SCM offline and do not specify one or more `scm_ranges`, the system selects which SCM increments to take offline.

`ONLINE` **or** `ON`
: The system brings the specified amount of SCM online. `ONLINE` is the default value if only `CONFIG SCM` is specified. The system rejects the command if you specify these values:

– A value that is not a multiple of the SCM increment size.
– A value that exceeds the total amount of SCM that is defined to this partition.
– A value that is not a valid amount of SCM (0, for example).
– SCM is not supported or not defined on the system.

`OFFLINE` **or** `OFF`
: The system takes the specified amount or specified ranges of SCM offline.

Attention: Taking SCM offline can affect data reliability and performance. Consider these implications before taking SCM offline:

- Your system must have enough auxiliary storage, which can include SCM and must include page data sets to back critical system data. The **`CONFIG SCM OFFLINE`** command fails if taking the specified number of SCMs offline results in leaving auxiliary storage more than 50% full.
- SCM is used for paging critical address spaces and common address spaces. An insufficient amount of SCM causes those address spaces to demote the larger pages to page-to-page data sets, which can lead to a loss of critical data during a direct access storage device (DASD) IBM HyperSwap® scenario.
- SCM is used for paging large pages. If there is an insufficient amount of SCM, 1 MB large pages are demoted to 256 4-KB pages and page-to-page data sets, which can negatively affect system performance.

`scm_ranges`
: Specifies a range of SCMs or a list of ranges that are separated by commas that are identified by *ddd{M|G|T}* - *ddd{M|G|T}*, for example, 0G - 16G or 32G - 64G. The starting and ending addresses for each range of SCMs must be multiples of the increment size.

The **`DISPLAY ASM`** and **`DISPLAY M`** commands are enhanced to display information and statuses that are related to VFM:

`DISPLAY ASM` Lists the SCM status along with the paging data set status.
`DISPLAY ASM,SCM` Displays a summary of SCM usage.
`DISPLAY M=SCM` Displays SCM online or offline and increment information.
`DISPLAY M=SCM(DETAIL)` Displays detailed increment-level information.

Tip: You might notice a difference in usage numbers between the **`DISPLAY M=SCM`** and **`DISPLAY ASM`** commands. The difference is because of how ASM perceives its usage of the cache of available SCM block IDs that ASM maintains. To ASM, some block IDs are not in use because they were not assigned to page-out requests. However, to the **`DISPLAY M=SCM`** command processor, block IDs are used because they were assigned to ASM for its use.

In 10.2.3, “Configuring VFM” on page 224, we allocate an initial VFM of 32 GB to the LPAR SPARE0, and a maximum VFM of 64 GB, as shown in Figure 10-40 on page 225. Now, from SPARE0, which is running the z/OS image HSSHCD, we run the IBM MVS **DISPLAY IPLINFO,PAGESCM** command. Example 10-3 shows the results.

Example 10-3 DISPLAY IPLINFO,PAGESCM

```
DISPLAY IPLINFO,PAGESCM
IEE255I SYSTEM PARAMETER 'PAGESCM': NOT_SPECIFIED
```

Because the **PAGESCM** parameter *is not* specified, the default value of `ALL` is used. If a VFM allocation is defined for the LPAR and `PAGESCM=ALL` is specified (or kept at the default), then at IPL time, the initial amount of VFM that is specified is used automatically by z/OS for paging. Likewise, if a specific amount is specified, this amount is made available for paging.

From HSSHCD, run the enhanced **DISPLAY ASM** and **DISPLAY M** commands to display the VFM SCM-related information and status. The result for each command is shown in Example 10-4.

Example 10-4 DISPLAY commands

```
DISPLAY ASM
IEE200I 11.54.04 DISPLAY ASM 259
TYPE     FULL STAT   DEV  DATASET NAME
PLPA      53%   OK  26E0  PAGE.HSSHCD.PLPA
COMMON     8%   OK  26E0  PAGE.HSSHCD.COMMON
LOCAL      0%   OK  26E0  PAGE.HSSHCD.LOCAL
SCM        0%   OK   N/A   N/A
PAGEDEL COMMAND IS NOT ACTIVE

DISPLAY ASM,SCM
IEE207I 11.55.26 DISPLAY ASM 261
STATUS      FULL              SIZE             USED         IN-ERROR
IN-USE        0%         8,388,608           14,307                0

DISPLAY M=SCM
IEE174I 11.56.19 DISPLAY M 263
STORAGE-CLASS MEMORY STATUS
64G DEFINED
ONLINE
0G-32G
32G OFFLINE-AVAILABLE
0% IN USE
SCM INCREMENT SIZE IS 16G

DISPLAY M=SCM(DETAIL)
IEE174I 11.57.11 DISPLAY M 272
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
64G DEFINED
ADDRESS  IN USE  STATUS
     0G      0%  ONLINE
    16G      0%  ONLINE
ONLINE: 32G  OFFLINE-AVAILABLE: 32G  PENDING OFFLINE: 0G
0% IN USE
SCM INCREMENT SIZE IS 16G
```

From these commands, you see that 64 GB of VFM is defined, but only 32 GB are online, and the other 32 GB are offline-available.

To vary an extra 32 GB VFM online to the example LPAR, run the **CONFIG SCM(xxG),ONLINE** command, as shown in Example 10-5. The amount of VFM that is configured online must be specified according to the supported increment size. From these displays, the supported increment size is `16G`.

Example 10-5 CONFIG SCM(16G),ONLINE

```
CONFIG SCM(32G),ONLINE
IEE195I SCM LOCATIONS 32G TO 64G ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Run the **DISPLAY ASM** and **DISPLAY M** commands again to display the status of the VFM and see that the 16 GB extra value is now online and available (Example 10-6).

Example 10-6 Post-configuration displays

```
DISPLAY ASM
IEE200I 12.01.43 DISPLAY ASM 284
TYPE     FULL STAT   DEV  DATASET NAME
PLPA      53%   OK  26E0  PAGE.HSSHCD.PLPA
COMMON     8%   OK  26E0  PAGE.HSSHCD.COMMON
LOCAL      0%   OK  26E0  PAGE.HSSHCD.LOCAL
SCM        0%   OK   N/A   N/A
PAGEDEL COMMAND IS NOT ACTIVE

DISPLAY ASM,SCM
IEE207I 12.02.36 DISPLAY ASM 286
STATUS      FULL              SIZE             USED         IN-ERROR
IN-USE        0%        16,777,216           14,307                0

DISPLAY M=SCM
IEE174I 12.03.27 DISPLAY M 288
STORAGE-CLASS MEMORY STATUS
64G DEFINED
ONLINE
 0G-64G
0G OFFLINE-AVAILABLE
0% IN USE
SCM INCREMENT SIZE IS 16G

DISPLAY M=SCM(DETAIL)
IEE174I 12.04.12 DISPLAY M 290
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
64G DEFINED
ADDRESS  IN USE  STATUS
     0G      0%  ONLINE
    16G      0%  ONLINE
    32G      0%  ONLINE
    48G      0%  ONLINE
ONLINE: 64G  OFFLINE-AVAILABLE: 0G  PENDING OFFLINE: 0G
0% IN USE
SCM INCREMENT SIZE IS 16G
```

When displaying the Storage Information windows on the SE again (compare to Figure 10-42 on page 226 and Figure 10-43 on page 227), this change in LPAR HSSHCD is reflected.

In Figure 10-44, the amount of allocated VFM went up to 144 GB (compared to the 112 GB in Figure 10-42 on page 226).

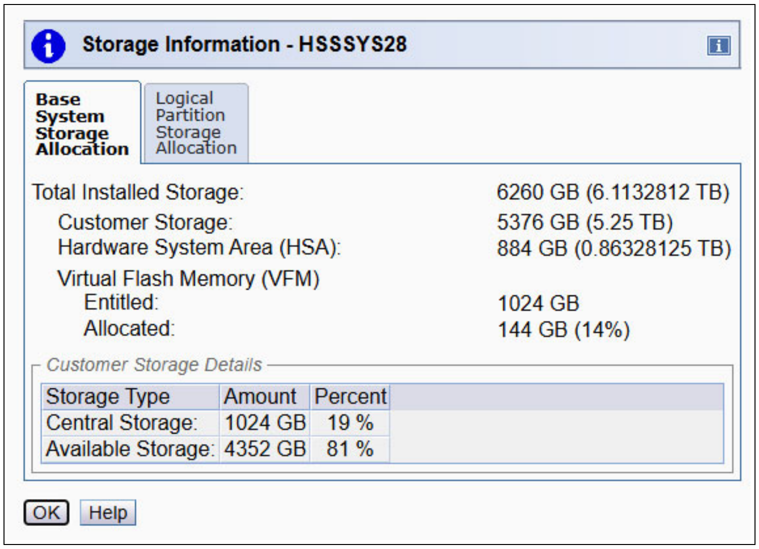

Figure 10-44 Results of CONFIG SCM(32G),ONLINE in LPAR SPAR0 - Base System Storage Allocation

In Figure 10-45, the amount of VFM allocated to LPAR SPARE0 went up to 64 GB.

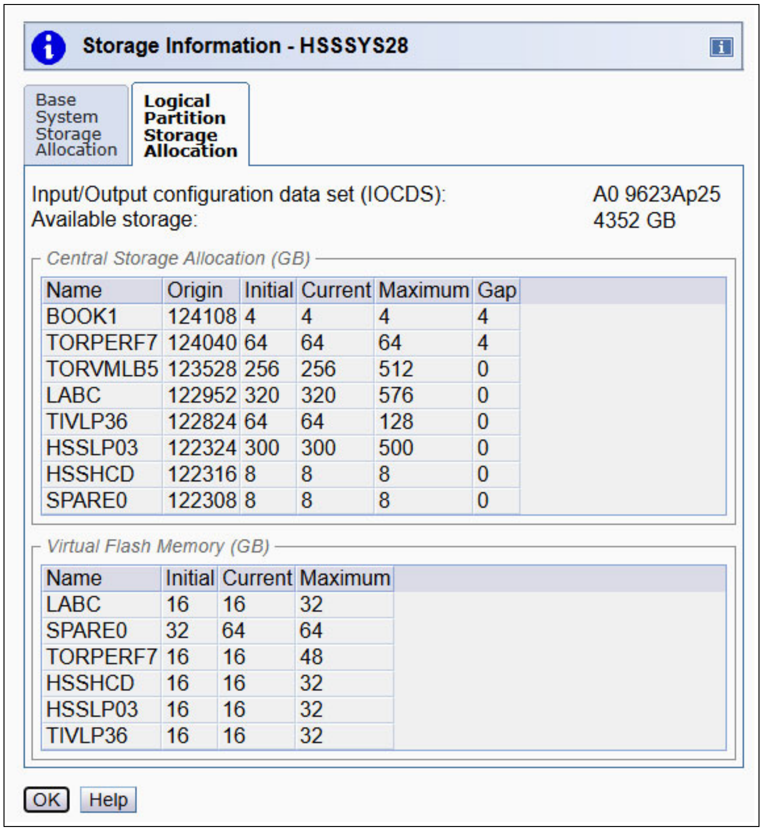

Figure 10-45 Results of CONFIG SCM(32G),ONLINE in LPAR SPARE0 - Logical Partition Storage Allocation

You can also set VFM offline, even to an amount that is lower than the initial value that is specified in the image activation profile. For LPAR SPARE0 where the amount of online VFM is reduced to 0 GB by running **CONFIG SCM(64G),OFFLINE**, you see the results in the Storage Information windows that are shown in Figure 10-46 and Figure 10-47. In Figure 10-46, the amount of allocated VFM went down to 80 GB.

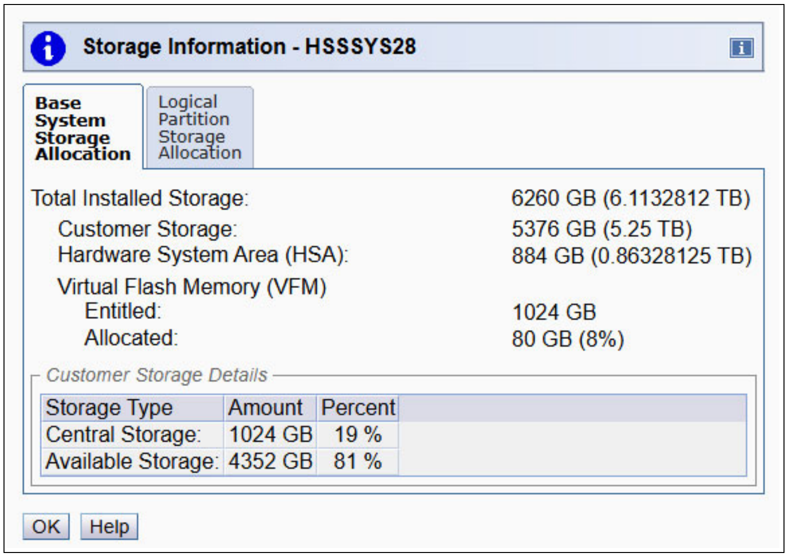

Figure 10-46 Results after CONFIG SCM(64G),OFFLINE in LPAR SPARE0 - Base System Storage Allocation

In Figure 10-47, the amount of VFM allocated to LPAR SPARE0 went down to 0 GB, which is lower than the initial 32 GB.

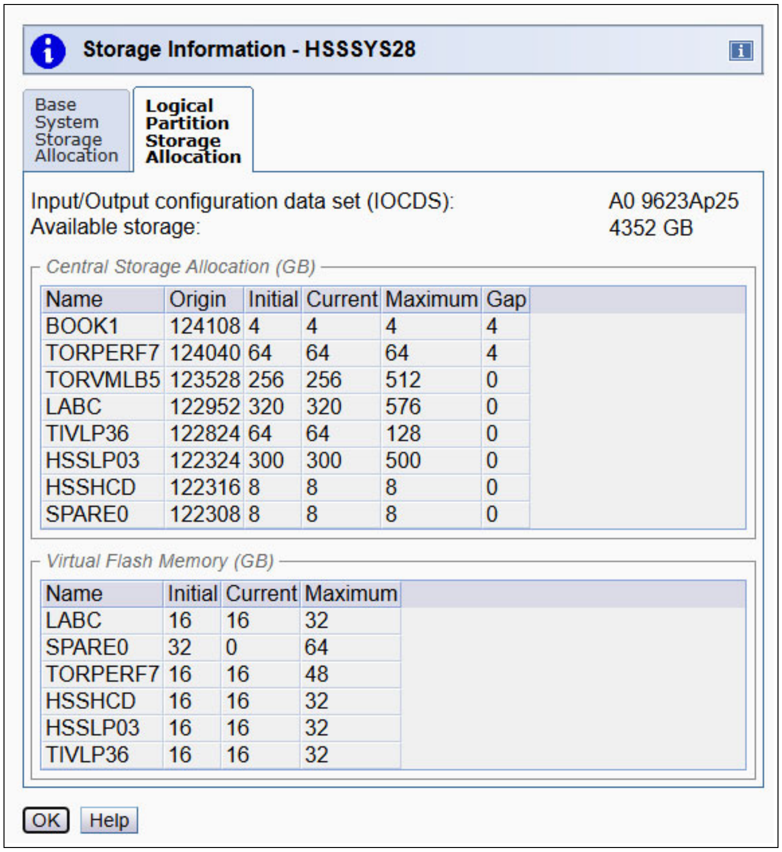

Figure 10-47 Results after CONFIG SCM(64G),OFFLINE in LPAR SPARE0 - Logical Partition Storage Allocation

Note: An LPAR uses only the amount of VFM activated for that LPAR. VFM that is set offline by the OS is returned to be used by other LPARs.

Important:

The allocation of VFM to a CF LPAR is no longer supported on z17 ME1. With CFLEVEL 26, which comes with the IBM z17 Driver Level 61, the support for use of Virtual Flash Memory for Coupling Facility images has been removed. If a client was using Virtual Flash Memory as an "overflow" mechanism for MQ Shared Queue structures, the client should plan to use other alternatives such as using larger Coupling Facility structures, using IBM MQ shared message data sets (SMDS), or using MQ techniques to offload MQ message data to Db2.

11

Adding logical partitions and operating system configurations

This chapter describes the steps for adding logical partitions (LPARs) and operating system configurations (OSCONFIGs) to your input/output definition file (IODF). It includes a list of potential configuration items and a short description of how to perform each task using a Hardware Configuration Definition (HCD).

Naming: This publication covers the IBM z17 ME1. Throughout this chapter, we refer to these machines as IBM z17.

Note: Not all configuration items are required for your installation. Additionally, the examples presented are not exhaustive.

This chapter focuses on the definition process. For a deeper understanding of how to use HCD, refer to these resources:

- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804
- Review *z/OS HCD Planning,* GA32-0907 in IBM Documentation

This chapter includes the following topics:

- Preliminary actions
- Defining an OSCONFIG
- Defining a logical partition (LPAR)

11.1 Preliminary actions

When defining new OSCONFIGs or LPARs, you might also need to define new devices such as consoles, network devices, and storage devices. These definition actions should be completed beforehand and are covered in the respective chapters of this book.

11.2 Defining a new OSCONFIG

An OSCONFIG describes the z/OS view of devices attached to a partition, including their operational behavior, parameters, and features. An LPAR is a hardware definition of a runtime environment for an operating system (OS) and the devices visible from it.

To create a new OSCONFIG, select **1,1** from the primary HCD menu. Enter **add** into the command line to bring up the panel shown in Figure 11-1. Then, enter the name of the new OSCONFIG and a description.

```
+---------------- Add Operating System Configuration -----------------+
|                                                                     |
|                                                                     |
| Specify or revise the following values.                             |
|                                                                     |
| OS configuration ID  . . . . . ITSOTEST                             |
| Operating system type  . . . . MVS      +                           |
|                                                                     |
| Description  . . . . . . . . . ITSO test OSCONFIG for IBM z17__     |
|                                                                     |
| OS config ID for D/R site  . . ________  (generated for GDPS)       |
|                                                                     |
|                                                                     |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap     |
| F12=Cancel                                                          |
+---------------------------------------------------------------------+
```

Figure 11-1 Add operating system configuration

When you return to the list of OS configurations, enter **S** before your new configuration to define a new Esoteric Device Table (EDT) for the OSCONFIG. Assign it a number, such as **00**, and provide a description as shown in Figure 11-2 on page 238.

```
+------------------------------- Add EDT --------------------------------+
|                                                                        |
|                                                                        |
| Specify the following values.                                          |
|                                                                        |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for IBM z17       |
|                                                                        |
| EDT identifier . . . 00                                                |
| Description  . . . . ________________________________                  |
|                                                                        |
|  F1=Help     F2=Split     F3=Exit     F5=Reset     F9=Swap     F12=Cancel |
+------------------------------------------------------------------------+
```

Figure 11-2 Add EDT

Next, enter **S** in front of the listed new EDT. On the list of esoterics, enter **add** into the command line to bring up the **Add Esoteric** panel. Add the required esoterics as shown in Figure 11-3.

```
+-------------------- Add Esoteric ---------------------+
|                                                        |
|                                                        |
| Specify the following values.                          |
|                                                        |
| Esoteric name  . . . ITSOES01                          |
| VIO eligible . . . . Yes    (Yes or No)                |
| Token  . . . . . . . 1___                              |
|                                                        |
|                                                        |
|  F1=Help    F2=Split   F3=Exit    F5=Reset   F9=Swap   |
| F12=Cancel                                             |
+--------------------------------------------------------+
```

Figure 11-3 Add Esoteric screen

Before adding devices to the esoterics, you must first add them to the OSCONFIG. Access the I/O Device List using the command **GOTO DEV**. Select **C** for the device or device group you want to add to the OSCONFIG. One of the subsequent panels will be **Define Device** to Operating System Configuration. Here, you can add the devices to your new OS configuration by entering **S** in front of the OSCONFIG.

After adding devices to the OSCONFIG, return to the Esoteric List of your OS configuration and add devices to the different esoterics.

11.3 Defining a new LPAR

An LPAR is a hardware definition of a runtime environment for an OS and the devices visible from it.

Defining a new LPAR consists of two separate actions:

1. Define the LPAR using HCD.
2. Define the Image profile on HMA.

To define an LPAR, you need at least the following information:

- Channel subsystem (CSS) where the LPAR is to be placed
- Partition name
- Partition number (identifier)
- Usage (for example, Coupling Facility)
- CHPIDs that will be accessible from this LPAR

11.3.1 Define LPAR using HCD

HCD initially defines a central processor complex (CPC) with all available channel subsystems (CSS) and LPAR predefined. The partitions are initially defined as reserved LPARs, meaning they are defined with a partition name of "*" and a usage of "CF/OS".

To add an LPAR to a CSS, first select **1.3 (Processors)** from the HCD main menu. Then, select the CPC where you want to add the LPAR. Next, select the list of partitions of the CSS where the new partition will be added by entering **P** in front of that CSS. This brings up the partition list as shown in Figure 11-4.

```
+----------------------------- Partition List -----------------------------+
|   Goto  Backup  Query  Help                                               |
| ------------------------------------------------------------------------ |
|                                                             Row 2 of 15 |
| Command ===> ________________________________________ Scroll ===> PAGE   |
|                                                                          |
| Select one or more partitions, then press Enter. To add, use F11.       |
|                                                                          |
| Processor ID . . . . : MENSA     Mensa                                  |
| Configuration mode  . : LPAR                                             |
| Channel Subsystem ID  : 3                                                |
|                                                                          |
| / Partition Name   Number Usage + UID Description                       |
| C *                2      CF/OS   N   ________________________________  |
| _ *                3      CF/OS   N   ________________________________  |
| _ *                4      CF/OS   N   ________________________________  |
| _ *                5      CF/OS   N   ________________________________  |
| _ *                6      CF/OS   N   ________________________________  |
| _ *                7      CF/OS   N   ________________________________  |
| _ *                8      CF/OS   N   ________________________________  |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset       |
|  F7=Backward   F8=Forward    F9=Swap      F10=Actions   F11=Add         |
| F12=Cancel    F13=Instruct  F22=Command                                  |
+--------------------------------------------------------------------------+
```

Figure 11-4 Partition List

You need to change the definition of a reserved LPAR by entering **C** in front of the chosen LPAR. This will bring up the **Change Partition Definition** panel, as shown in Figure 11-5.

```
+-------------- Change Partition Definition ---------------+
|                                                           |
|                                                           |
| Specify or revise the following values.                  |
|                                                           |
| Partition name . . . MENSA32                              |
| Partition number . . 2       (same as MIF image ID)       |
| Partition usage  . . OS      +                            |
| UID uniqueness . . . N        (Y/N)                       |
|                                                           |
| Description  . . . . MENSA32 test OS partition           |
|                                                           |
|                                                           |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset     |
|  F9=Swap   F12=Cancel                                     |
+-----------------------------------------------------------+
```

Figure 11-5 Change Partition Definition

Change the partition name to the desired name and update the usage to CF for a Coupling LPAR or OS for an LPAR that will run an OS. Additionally, add a description for the partition. Do not change the partition number.

If you plan to rename a partition, HCD allows you to do this with a single change of the partition definition. However, activating an I/O definition containing such a change will fail because you cannot rename existing, non-reserved partitions. Therefore, you must follow these three steps:

1. Define an I/O configuration with the partition changed to a reserved partition, stop and deactivate the LPAR, and activate the new configuration.
2. On the HMA, change the name of the image profile to the new partition name by opening it with **Customize profile**, changing the name at the top of the definition panel, and selecting **Save**.
3. Define another I/O configuration starting from the original I/O definition with the partition changed to the desired name, and activate the new configuration.

Using the original I/O definition as the source for the third step helps retain the original CHPIDs attached to the partition. Otherwise, this information would be lost after changing the partition to reserved and back to a named partition.

11.3.2 Define image profile on HMA

After activating the I/O configuration that contains the new LPAR definition, define the image profile on the HMA as shown in 5.4, “Creating an image profile on the Support Element” on page 105.

12 Adding storage devices

This chapter describes how to define Fibre Connection (FICON) switches and FICON channel path IDs (CHPIDs), and connect them to a direct access storage device (DASD) control unit (CU). It includes a list of these potential configuration items and a short description about how to configure each of them by using Hardware Configuration Definition (HCD).

This chapter shows only the definition process. If you want a deeper understanding about how to use HCD and FICON, see these resources:

- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804
- *FICON Planning and Implementation Guide*, SG24-6497
- *z/OS HCD Planning,* GA32-0907

Naming: The IBM z17 systems that are targeted by this publication consist of IBM z17 ME1. Throughout this chapter, we might refer to these machines as IBM z17.

Note: The examples that are shown in this chapter are based on the IBM z17 ME1 (9175).

Note: Not all the following configuration items are necessarily required for your installation. In addition, the examples that are presented are not exhaustive.

This chapter includes the following topics:

- Defining more I/O by using HCD
- FICON CHPIDs, switches, and DASD CUs

12.1 Defining more I/O by using HCD

When defining new I/O components in an input/output definition file (IODF), certain definitions like operating system configurations (OSCONFIGs), partitions, FICON switches, CUs, and devices must be done first. After these items are defined, the connections can be made:

- Defining FICON switches (directors, storage area networks, and storage area network switches)
- Defining FICON CHPIDs and connecting them to a FICON switch
- Defining FICON CHPIDs for a direct connection to a 2107 control unit
- Defining FICON CHPIDs for switch connections to a 2107 control unit
- Defining 3390B devices to an OSCONFIG and Eligible Device Tables and Esoteric device groups

The following I/O definitions use HCD to demonstrate the examples. The examples continue by using the work example `SYS9.IODF81.WORK`.

12.1.1 Defining FICON switches (directors, storage area networks, and storage area network switches)

The following items are considerations for a new FICON switch:

- Switch ID
- Switch type
- Serial number (optional)
- Description (optional)
- Switch address
- Installed port range
- Switch CU number
- Switch device number

To add a FICON Switch by using HCD, complete the following steps:

1. From the main HCD panel, select option `1.2. Switches`.
2. In the CLI, enter `add` (see Figure 12-1 on page 245) to add a switch.
3. Make the following updates (we used a switch ID value of `41`), and press Enter:
   - Update `Switch ID` to `41`.
   - Update `Switch type` to `2032`.
   - Update `Description` to `ITSO test storage area network definition`.
   - Update `Switch address` to `41`.
   - Update `Installed port range` to `00` and `FE`.
   - Update `Switch CU number` to `0041`.
   - Update `Switch device number` to `0041`.

```
+---------------------------- Add Switch ----------------------------+
|                                                                    |
|                                                                    |
| Specify or revise the following values.                            |
|                                                                    |
| Switch ID . . . . . . . . 41  (00-FF)                              |
| Switch type . . . . . . . 2032           +                         |
| Serial number . . . . . . __________                               |
| Description . . . . . . . ITSO test SAN definition______           |
| Switch address  . . . . . 41  (01-EF) for a FICON switch           |
|                                                                    |
| Specify the port range to be installed only if a larger range      |
| than the minimum is desired.                                       |
|                                                                    |
| Installed port range  . . 00  -  FE  +                             |
|                                                                    |
| Specify either numbers of existing control unit and device, or     |
| numbers for new control unit and device to be added.               |
|                                                                    |
| Switch CU number(s) . . . 0041  ____  ____  ____  ____  +          |
| Switch device number(s) . 0041  ____  ____  ____  ____             |
|                                                                    |
|                                                                    |
+--------------------------------------------------------------------+
```

Figure 12-1 Switches: Add Switch

Figure 12-2 shows the new FICON switch definition.

```
 Switch List          Row 1 of 3 More:       >
 Command ===> _______________________________________________ Scroll ===> CSR

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

                                                                   CU   Dev
 / ID Type +        Ad Serial-# + Description                      Num. Num.
 _ 01 2032          01 10546MH    8960-F64 SAN64B-6 SW 01          0001 0001
 _ 02 2032          02 10546MD    8960-F64 SAN64B-6 SW 03          0002 0002
 _ 41 2032          41 __________ ITSO test SAN definition         0041 0041
 ****************************** Bottom of data *******************************







+---------------------------------------------------------------------+
| Switch control unit(s) 0041 and device(s) 0041 defined, but not yet |
| connected to both a processor and an operating system.              |
+---------------------------------------------------------------------+
```

Figure 12-2 Switches: FICON switch added

12.2 FICON CHPIDs, switches, and DASD CUs

This section describes the following topics:

- Defining FICON CHPIDs and connecting them to a FICON switch
- Defining FICON CHPIDs for a direct connection to a 2107 control unit
- Defining FICON CHPIDs for switch connections to a 2107 control unit
- Defining 3390B devices to an OSCONFIG and Eligible Device Tables and Esoteric device groups

12.2.1 Defining FICON CHPIDs and connecting them to a FICON switch

Here are considerations for a new FICON CHPID:

- CHPID
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Dynamic entry switch ID
- Entry switch ID
- Entry port
- Partition access list

Here are more considerations:

- For performance and redundancy, determine how many CHPIDs are required to connect to the FICON switch and then to one or more CUs.
- For performance and redundancy, determine how many I/O cards of that feature are installed in the processor and to what Peripheral Component Interconnect Express (PCIe) ports and on which central processor complex (CPC) drawer does the I/O cards connect to. (For a list of installed hardware, see the physical channel ID (PCHID)/CHID report.)
- After a CHPID is defined, it can be added to a predefined partition in that channel subsystem (CSS).
- After a CHPID is defined, it can be connected to a FICON switch.
- After a CHPID is defined, it can be connected to a CU (covered in a later step).

To define a FICON CHPID, connect to a FICON switch and provide access to a partition. Then, complete the following steps[1]:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. Press PF11 or in the CLI, enter `add` (see Figure 12-3 on page 247) to add a CHPID.
5. Make the following updates, and press Enter:
   - Update `Channel path ID` to `14`.
   - Update `Channel ID` to `114`.
   - Update `Channel path type` to `FC`.
   - Update `Operational mode` to `SHR`.

[1] in this book, we previously recommended to leave the PCHID blank and assign it using the Connectivity Mapping Tool (CMT). For completeness. these steps show how to assign them by using the HCD.

- Update `Description` to whatever description that you want.
- Update `Dynamic entry switch ID` to `41`.
- Update `Entry switch ID` to `41` (optional but preferred).
- Update `Entry port` to `10` (optional but preferred).

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . 14    +            Channel ID  114  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . FC_   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0388 32Gb FICON Exp32G SX____                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  41  + (00 - FF)                                  |
| Entry switch ID  . . . . 41  +                                            |
| Entry port . . . . . . . 10  +                                            |
|                                                                           |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 12-3 Processors: Add Channel Path

HCD now prompts you to select which partition the CHPID should have access to.

6. Type forward slash (/) next to the partition you want (see Figure 12-4), and press Enter.

```
+--------------------------- Define Access List ----------------------------+
|                                                                Row 1 of 1 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.           |
|                                                                           |
| Channel subsystem ID : 3                                                  |
| Channel path ID  . . : 14     Channel path type  . : FC                   |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 3      MENSA32          2      OS    MENSA32 test OS partition          |
| *************************** Bottom of data ****************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 12-4 Processors: Define Access List

Because we have only one partition that is defined in this CSS, we do not see the Define Candidate List panel, even though we defined the CHPID as SHR.

HCD returns to the Channel Path List and shows the CHPID that was defined (see Figure 12-5).

```
Channel Path List       Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 14    114   FC    SHR   41  41 10       No  FC#0388 32Gb FICON Exp32G SX
```

Figure 12-5 Processors: Channel Path List

12.2.2 Defining FICON CHPIDs for a direct connection to a 2107 control unit

Here are considerations for a new FICON CHPID:

- CHPID
- CHID
- Channel path type
- Operational mode
- Description
- Partition access list

The direct connect method is used in an environment where there is only one processor. The FICON switch method, which is described in 12.2.3, “Defining FICON CHPIDs for switch connections to a 2107 control unit” on page 252, is used where multiple processors must connect to the same CUs. This situation might not always be the case.

Here are considerations for connecting a FICON CHPID to a DASD CU and its devices:

- For performance and redundancy, determine how many I/O cards of that feature are installed in the processor and to what PCIe ports on what CPC drawer does the I/O cards connect to. (For a list of installed hardware, see the PCHID/CHID report.)
- After a CHPID is defined, it can be added to a predefined partition in that CSS.
- The CHPID can be connected to a CU.

> **Note:** For FICON Express16SA (Feature Code #0436 and Feature Code #0437), FICON Express32S (Feature Code #0461 and Feature Code #0462), and FICON Express32G (Feature Code #0387 and Feature Code #0388), defining both Fibre Channel (FC) and Fibre Channel Protocol (FCP) CHPID types on the same I/O card is not supported. (There is no mix that is supported by IBM z16 and IBM z17). HCD issues the following error message during a Validate or Build Production for an IODF:
>
> `CBDA964I Chpid type mix detected on processor MENSA for channels: 1.6C, 0.6D`

To define a FICON CHPID that connects directly to a CU and provides access to a partition, complete the following steps[2]:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. Press PF11 or in the CLI, enter `add` to add a CHPID.

[2] in this book, we previously recommended to leave the PCHID blank and assign it using the CMT. For completeness. these steps show how to assign them using the HCD.

5. Make the following updates (see Figure 12-6), and press Enter:
   - Update `Channel path ID` to `15`.
   - Update `Channel ID` to `115`.
   - Update `Channel path type` to `FC`.
   - Update `Operational mode` to `SHR`.
   - Update `Description` to the description that you want.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . 15    +            Channel ID  115  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . FC_   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0388 32Gb FICON Exp32G SX____                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 12-6 Processors: Add Channel Path

6. Type forward slash (/) next to the partition that you want CHPID to have access to (see Figure 12-7 on page 251), and press Enter.

```
+--------------------------- Define Access List ----------------------------+
|                                                                Row 1 of 1 |
| Command ===> ________________________________________ Scroll ===> CSR     |
|                                                                           |
| Select one or more partitions for inclusion in the access list.          |
|                                                                           |
| Channel subsystem ID : 3                                                  |
| Channel path ID  . . : 15     Channel path type  . : FC                   |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 3      MENSA32          2      OS    MENSA32 test OS partition          |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 12-7 Processors: Define Access List

Because only one partition is defined in this CSS, we do not see the Define Candidate List panel, even though we defined the CHPID as SHR.

HCD now returns to the Channel Path List panel showing the CHPID that was defined (see Figure 12-8).

```
 Channel Path List       Row 1 of 2 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 14    114   FC    SHR   41  41 10       No  FC#0388 32Gb FICON Exp32G SX
_ 15    115   FC    SHR   __  __ __       No  FC#0388 32Gb FICON Exp32G SX
******************************* Bottom of data ********************************
```

Figure 12-8 Processors: Channel Path List

12.2.3 Defining FICON CHPIDs for switch connections to a 2107 control unit

There are two ways to define FICON connections to a CU: One is through a FICON switch, and the other is direct connect, which is described in 12.2.2, “Defining FICON CHPIDs for a direct connection to a 2107 control unit” on page 249.

The direct connect method is used in an environment where there is only one processor. The FICON switch method is used where multiple processors must connect to the same CUs. This situation might not always be the case.

- For FICON switch connections, there is usually a minimum of two FICON switches that the FICON CHPIDs connect through, primarily for failure or service redundancy of the FICON switches.

> **Note:** The device type of the CU depends on your storage server. An IBM DS8000 server uses a CU type of 2107.

- For this example, we connect to a predefined 2107 CU (A000), with a control unit address (CUADD) of 40 and devices A000-A0EF (3390B) and A0F0-A0FF (3390A).

> **Note:** For FICON Express16SA (Feature Code #0436 and Feature Code #0437), and FICON Express32S (Feature Code #0461 and Feature Code #0462), and FICON Express32G (Feature Code #0387 and Feature Code #0388), defining both FC and FCP CHPID types on the same I/O card is not supported. (There is no mix that is supported by IBM z16 and IBM z17). HCD issues the following error message during a Validate or Build Production for an IODF:
>
> `CBDA964I Chpid type mix detected on processor MENSA for channels: 1.6C, 0.6D`

To define FICON CHPIDs for a switch connection, complete the following steps:

1. From the main HCD panel, select option `1.4. Control units`.
2. Scroll through the CU list until you find the CU that you want to connect to, or in the CLI, enter `L A000`. In our example, we use `A000`.
3. Enter `c` next to the CU definition, and press Enter.

4. Make the following updates for a FICON switch connection (see Figure 12-9), and press Enter:
   a. Update `Connected to switches` to `41`. Switch 41 is the switch ID that we defined in the previous example.
   b. Update `Ports` to `11`. Port 11 is the FICON switch port that the CU connects to (also known as the Destination Link Address (DLA) of the FICON CHPID).

```
+----------------- Change Control Unit Definition ------------------+
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . A000  +                              |
|                                                                   |
| Control unit type  . . . . . 2107          +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . DS8950F - SU2 - SSID=A000            |
|                                                                   |
| Connected to switches  . . . 41 __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . 11 __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
+-------------------------------------------------------------------+
```

Figure 12-9 Control units: Change Control Unit Definition

HCD now shows the Select Processor / CU panel. This panel is where the connection is now made between the FICON CHPID (14) and the CU FICON Switch Port (11) through the FICON switch (41).

5. Enter `c` next to the `Processor.CSS` that contains the partition that we want to have access to the CU and also has access to the CHPID that we want to connect to the CU. In our example, we use `MENSA.3`, and then press Enter.

6. Make the following updates to define the Processor / CU connection panel (see Figure 12-10), and press Enter:
   a. Update `Channel path IDs` to `14`. CHPID 14 is the CHPID that we defined in the previous example.
   b. Update `Link address` to `4111`. Link address 4111 is FICON switch 41 and CU Port 11.
   c. Update `Unit address` to `00`. Unit Address (UA) of 00 is the starting UA number on the CHPID.
   d. Update `Number of units` to `256`. The number of units of 256 equals A000 - A0FF (that is, 00 - FF or 256).
   e. Update `Logical address` to `40`. The logical address of 40 is the CUADD that is defined in the DS8950F that defines the location of the devices in the DS8950F.

```
                              Select Processor / CU
+---------------------- Change Control Unit Definition -----------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : A000        Type . . . . . . : 2107                |
| Processor ID . . . . . : MENSA       Mensa                                  |
| Channel Subsystem ID . : 3                                                  |
|                                                                             |
| Channel path IDs . . . . 14    __   __   __   __   __   __   __    +        |
| Link address . . . . . . 4111  ____ ____ ____ ____ ____ ____ ____  +        |
|                                                                             |
| Unit address . . . . . . 00    __   __   __   __   __   __   __    +        |
| Number of units  . . . . 256   ___  ___  ___  ___  ___  ___  ___            |
|                                                                             |
| Logical address  . . . . 40  + (same as CUADD)                              |
|                                                                             |
| Protocol . . . . . . . . __  + (D, S or S4)                                 |
| I/O concurrency level  . _   + (1, 2 or 3)                                  |
|                                                                             |
|                                                                             |
|                                                                             |
+-----------------------------------------------------------------------------+
```

Figure 12-10 Control units: Change Control Unit Definition detail

HCD now shows the Modify Device Parameters panel where you can override the UA numbers. For a 2107 DASD CU definition, the starting UA is usually `00` and the ending UA is `FF`, giving you 256 DASD definitions for the CU.

7. In our example, we do not change the defaults that are proposed by HCD (see Figure 12-11). Press Enter.

```
+------------------------- Modify Device Parameters -------------------------+
|                                                     Row 1 of 2 More:      > |
| Command ===> __________________________________________ Scroll ===> PAGE   |
|                                                                             |
| Specify or revise any changes to the device parameters in the list below.  |
| To view attached control units, scroll to the right.                       |
|                                                                             |
| Processor ID . . . . : MENSA        Mensa                                  |
| Channel Subsystem ID : 3                                                    |
|                                                                             |
| ---------Device---------     ---UA----                Preferred Exposure   |
| No., Range Type         SS+ Old New + Time-Out STADET CHPID +   Device     |
| A000,240   3390B        _   00  00    No       Yes    __                   |
| A0F0,016   3390A        1   F0  F0    No       Yes    __                   |
| **************************** Bottom of data ***************************** |
|                                                                             |
+-----------------------------------------------------------------------------+
```

Figure 12-11 Control units: Modify Device Parameters

HCD returns to the Select Processor / CU panel, which shows the CHPID (14) to DLA (11) through the FICON switch (41) connection definition (see Figure 12-12).

```
                              Select Processor / CU    Row 1 of 15 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

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

Control unit number . . : A000     Control unit type . . . : 2107

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ MENSA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.1    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.3    14.4111 _______ _______ _______ _______ _______ _______ _______
_ MENSA.4    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.5    _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.0     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.1     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.2     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.3     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.4     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.5     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.0     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.1     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.2     _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 12-12 Processors: Select Processor / CU: CHPID to Link address connection

By pressing F20 (Right), you can see the other parts of the definition summary.

For a FICON direct connection to this CU definition, you omit the Link address values.

To define the Processor / CU connection (see Figure 12-13), update `Channel path IDs` to `15`, and press Enter.

```
                               Select Processor / CU
+---------------------- Change Control Unit Definition -----------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : A000        Type . . . . . . : 2107                |
| Processor ID . . . . . : MENSA       Mensa                                  |
| Channel Subsystem ID . : 3                                                  |
|                                                                             |
| Channel path IDs . . . . 14    15    __    __    __    __    __    __    +  |
| Link address . . . . . . 4111  ____  ____  ____  ____  ____  ____  ____  +  |
|                                                                             |
| Unit address . . . . . . 00    __    __    __    __    __    __    __    +  |
| Number of units  . . . . 256   ___   ___   ___   ___   ___   ___   ___      |
|                                                                             |
| Logical address  . . . . 40  + (same as CUADD)                              |
|                                                                             |
| Protocol . . . . . . . . __  + (D, S or S4)                                 |
| I/O concurrency level  . _   + (1, 2 or 3)                                  |
|                                                                             |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap        |
| F12=Cancel                                                                  |
+-----------------------------------------------------------------------------+
```

Figure 12-13 Control units: Change Control Unit Definition detail

8. HCD again shows the Modify Device Parameters panel where you can override the UA numbers. Press Enter to continue.

   HCD returns to the Select Processor / CU panel showing the CHPID (15) connection definition (see Figure 12-14 on page 258).

```
                              Select Processor / CU    Row 1 of 15 More:       >
Command ===> ______________________________________________ Scroll ===> PAGE

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

Control unit number . . : A000     Control unit type . . . : 2107

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ MENSA.3    14.4111 15      _______ _______ _______ _______ _______ _______
_ MENSA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.1    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.4    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.5    _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.0     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.1     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.2     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.3     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.4     _______ _______ _______ _______ _______ _______ _______ _______
_ PAVO.5     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.0     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.1     _______ _______ _______ _______ _______ _______ _______ _______
_ VELA.2     _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 12-14 Processors: Select Processor / CU: CHPID to Link address connection

Although a mixture of FICON switched and FICON direct connections are not recommended to the same CU, this configuration is possible.

A typical scenario might be where you were moving from direct connect DASD to a FICON switch connected DASD, but you cannot take the DASD offline to live systems.

12.2.4 Defining 3390B devices to an OSCONFIG and Eligible Device Tables and Esoteric device groups

The OSCONFIG name is the part of an IODF that determines what devices a z/OS system may access when it starts. Also, the partition that the z/OS system is restarted in also must have access to the CHPIDs that connect to the CUs and devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups that are defined in Eligible Device Tables (EDTs) within an OSCONFIG.

Esoterics device groups are used to request allocation of a device that was defined in an Esoteric device group when using the **`UNIT =`** parameter in a **`JCL DD`** statement. However, this allocation can be overridden or intercepted by using DFSMS.

The OSCONFIG name includes these items:

- Device Parameters and Features
- EDT ID: Esoterics / VIO
- NIP Consoles

Here are the considerations for adding devices to an OSCONFIG and Esoteric:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system has access to that device.
- Does the device need to be added to an Esoteric device group? Mostly, this item is installation-specific.
- The example adds the predefined 2107 devices A000 - A0EF (3390B) and A0F0 - A0FF (3390A).

To define 3390B devices to an OSCONFIG and EDT/Esoteric, complete the following steps:

1. From the main HCD panel, select option `1.5. I/O Devices`.
2. Scroll through the I/O Device List until you find the device number that you want to add to the OSCONFIG, or in the CLI, enter `L A000`. In our example, we use `A000`.
3. Enter `c` next to one or more device numbers, and press Enter.

   HCD displays the Change Device Definition panel where you can modify the CU that the devices are attached to (see Figure 12-15). Press Enter.

```
+------------------------- Change Device Definition --------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number  . . . . . . . : A000  (0000 - FFFF)                          |
| Number of devices  . . . . . : 240                                          |
| Device type  . . . . . . . . : 3390B                                        |
|                                                                             |
| Serial number  . . . . . . . . __________  +                                |
| Description  . . . . . . . . . ________________________________             |
|                                                                             |
| Volume serial number . . . . . ______  + (for DASD)                         |
|                                                                             |
| PPRC usage . . . . . . . . . . _  + (for DASD)                              |
|                                                                             |
| Connected to CUs . A000  ____  ____  ____  ____  ____  ____  ____  +        |
|                                                                             |
| ENTER to continue.                                                          |
|                                                                             |
|                                                                             |
|                                                                             |
+-----------------------------------------------------------------------------+
```

Figure 12-15 I/O Devices: Change Device Definition

HCD displays the Device / Processor Definition panel where you can modify some of the Device parameters relating to Subchannel Set (SS), UA, and the Explicit Device Candidate List.

4. Enter `c` next to the `Proc.CCSID` item if you want to modify any of these parameters (see Figure 12-16), or press Enter to accept the defaults.

```
+---------------------- Device / Processor Definition -----------------------+
|                                                                  Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> PAGE   |
|                                                                              |
| Select processors to change device/processor definitions, then press        |
| Enter.                                                                       |
|                                                                              |
| Device number  . . : A000       Number of devices  . : 240                   |
| Device type  . . . : 3390B                                                   |
|                                                                              |
|                                          Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null  |
| c MENSA.3     0    00   No        Yes     __         No             ___   |
| ***************************** Bottom of data ***************************** |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
+-----------------------------------------------------------------------------+
```

Figure 12-16 I/O Devices: Device / Processor Definition continued

5. Next is the HCD panel, where you define devices to the OSCONFIG. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG that you want to add to the devices to, or in the CLI, enter `L ITSOTEST`. In our example, we use `ITSOTEST`.
6. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the device parameters and features that are applicable to that device type. In our example, we add 3390B and 3390A devices to `ITSOTEST`.

7. Make the following updates to define the Processor / CU connection (see Figure 12-17), and press Enter:
   – Update `OFFLINE` to `No` (if you want these devices to be online during IPL time).
   – Update `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   – Update `LOCANY` to `Yes` (if the device UCB can be in 31-bit storage).
   – Update `WLMPAV` to `Yes` (if you want the device to be managed by Workload Manager).
   – Update `READ-ONLY` to `No` (use to set secondary devices to read only).
   – Update `SHARED` to `Yes` (if the device is going to be shared between other systems).
   – Update `SHAREDUP` to `No` (must be set to `No` if Shared is set to `Yes`).

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                 Row 1 of 7 |
| Command ===> ___________________________________________ Scroll ===> PAGE  |
|                                                                            |
| Specify or revise the values below.                                        |
|                                                                            |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for IBM z17           |
| Device number  . . : A000         Number of devices  : 240                 |
| Device type  . . . : 3390B                                                 |
|                                                                            |
| Parameter/                                                                 |
| Feature    Value +         R Description                                   |
| OFFLINE    No                Device considered online or offline at IPL    |
| DYNAMIC    Yes               Device supports dynamic configuration         |
| LOCANY     Yes               UCB can reside in 31 bit storage              |
| WLMPAV     Yes               Device supports work load manager             |
| READ-ONLY  No                Restrict access to read requests (SEC or NO)  |
| SHARED     Yes               Device shared with other systems              |
| SHAREDUP   No                Shared when system physically partitioned     |
| **************************** Bottom of data ****************************** |
|                                                                            |
|                                                                            |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 12-17 I/O Devices: Define Device Parameters / Features

The Assign/Unassign Device to Esoteric panel opens, where you can specify which Esoteric (if any) that you want the devices to be added to.

8. In the Assign/Unassign the Device to Esoteric panel (see Figure 12-18) under `Assigned`, enter `Yes`, and press Enter.

```
+-------------------- Assign/Unassign Device to +-oteric --------------------+
|                                                                  Row 1 of 1 |
| Command ===> ________________________________________ Scroll ===> PAGE      |
|                                                                              |
| Specify Yes to assign or No to unassign.  To view devices already           |
| assigned to esoteric, select and press Enter.                               |
|                                                                              |
| Configuration ID : ITSOTEST         ITSO test OSCONFIG for IBM z17          |
| Device number  . : A000             Number of devices  : 240                |
| Device type  . . : 3390B            Generic  . . . . . : 3390               |
|                                                                              |
| / EDT.Esoteric  Assigned  Starting Number Number of Devices                 |
| _ 00.ITSOES01   Yes       ____            ____                              |
| ***************************** Bottom of data ***************************** |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
|                                                                              |
+-----------------------------------------------------------------------------+
```

Figure 12-18 I/O Devices: Assign/Unassign Device to Esoteric

9. The final panel opens and shows that the devices are defined to the OSCONFIG (see Figure 12-19). Press Enter to return to the I/O Device List.

```
+----------- Define Device to Operating System Configuration -----------+
|                                                              Row 1 of 11 |
| Command ===> _____________________________________ Scroll ===> CSR   |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : A000          Number of devices  : 240             |
| Device type  . . : 3390B                                              |
|                                                                       |
| Subchannel Sets used in processor configurations  : 0                 |
|                                                                       |
| / Config. ID   Type     SS Description                     Defined    |
| _ ITSO         MVS         All ITSO devices                            |
| _ ITSOBLD      MVS         ITSO + CPO devices                          |
| _ ITSOTEST     MVS      0  ITSO test OSCONFIG for IBM z17   Yes       |
| _ PRF1         MVS         Proof Points                                |
| _ WAT1         MVS         Watsonx                                     |
| _ ZOSMAINT     MVS         z/OS Central Mgmt Image                     |
| _ ZOSMESUR     MVS         z/OS systems to monitor all                 |
| _ ZOS24BAS     MVS         z/OS 2.4 base system                        |
| _ ZOS24P1K     MVS         z/OS 2.4 P1K study system                   |
| _ ZOS25BAS     MVS         z/OS 2.5 inital system                      |
| _ ZOS25POP     MVS         z/OS 2.5 POP system                         |
+-----------------------------------------------------------------------+
```

Figure 12-19 I/O Devices: Assign/Unassign Device to Esoteric completed

The same steps can now be performed for the remaining devices A0F0 - A0FF (3390A) for this example.

13

Adding network devices

This chapter describes how to define OSC (Integrated Console Controller), OSD (Queued Direct I/O - QDIO), OSH (Enhanced QDIO), and Internal Queued Direct (IQD) channel path IDs (CHPIDs), along with the respective control units (CUs) and devices.

For more information, see the *Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7183.

> **Note:** IBM z17 does not support **CHPID type OSE**.

A Channel path is the connection between the channel subsystem and control units that allows the channel subsystem to communicate with I/O devices.

A control unit provides the capability necessary to operate and control an I/O device.

An I/O device is the end of the communication link.

In this chapter we provide a list of these potential configuration items and a short description about how to do each of them by using Hardware Configuration Definition (HCD).

> **Naming:** The IBM z17 system that are targeted by this publication consist of IBM z17 ME1(9175). Throughout this chapter, we might refer to this machine as IBM z17.
>
> **Note:** Not all the following configuration items are necessarily required for your installation. In addition, the examples that are presented are not exhaustive.

This chapter includes the following topics:

- Summary of network devices defined in this chapter
- OSC definitions (Integrated Console Controller)
- OSD definitions (Queued Direct I/O - QDIO)
- Network Express for channel type OSH
- IQD CHPIDs for HiperSockets

13.1 Summary of network devices defined in this chapter

When defining I/O components in an input/output definition file (IODF), certain definitions like operating system configurations (OSCONFIGs), partitions, Fibre Connection (FICON) switches, CUs, and devices must be done first. After these items are defined, then the additional definitions can be made:

- OSC definitions (Integrated Console Controller)
- OSD definitions (Queued Direct I/O - QDIO)
- Network Express for channel type OSH
- IQD CHPIDs for HiperSockets

Note: We created our definitions on a z/OS Version 3, Release 1. All definitions and values are examples only.

13.2 OSC definitions (Integrated Console Controller)

This section covers defining OSC CHPID, control unit and devices.

13.2.1 Defining OSC CHPIDs

The OSA-ICC functionality requires an unique CHPID, control unit and device definition. These definitions are made using HCD or IOCP, which defines the OSA-Express feature to the I/O hardware configuration.

When defining an OSC connection, first you must determine which type of Open Systems Adapter-Express (OSA-Express) feature that you need for your configuration. See Table 13-1.

Table 13-1 OSA Express required to channel type OSC.

Networking Features	Feature Code (FC)	Ports per FC
OSA Express7S 1.2 SX 1G	0455	2
OSA Express7S 1.2 LX 1G	0454	2

Note: Starting on IBM z17, OSA Express7S 1000BaseT will only support channel path type OSD.

Here are the parameters for the definition of OSC CHPID:

- Channel path id (CHPID).
- Channel ID (CHID).
- Channel path type.
- Operational mode.
- Description.
- Partition access list.
- For performance and redundancy, determine how many I/O cards of that feature are installed in the processor and to what Peripheral Component Interconnect Express (PCIe) ports on what central processor complex (CPC) drawer that the I/O cards connect to. (For a list of installed hardware, see the physical channel ID (PCHID)/CHID report.)

To define an OSC CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the channel subsystem (CSS) ID that you want to add a CHPID to, and press Enter.
4. Press PF11 or, in the command line, enter `add` (see Figure 13-1) to add a CHPID.
5. Make the following updates, and press Enter:
   a. Update `Channel path ID` to `B1`.
   b. Update `Channel ID` to `1D4`.
   c. Update `Channel path type` to `OSC` (To define the OSA-ICC function).
   d. Update `Operational mode` to `SHR` (To share this channel path among logical partitions or even SPAN to share over multiple CSSs).
   e. Update `Description` to the description that you want.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . B1    +           Channel ID  1D4  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSC   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0454 OSA-Express7S 1.2 GbE LX                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
+---------------------------------------------------------------------------+
```

Figure 13-1 Processors: Add Channel Path - OSC

Note: We recommend using a meaningful description, which will serve as a reference point in HCD.

6. HCD prompts you to select the partitions that will have access to the CHPID. Enter a forward slash (/) next to the partitions that you want (see Figure 13-2) and press Enter.

```
+--------------------------- Define Access List ----------------------------+
|                                                              Row 12 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.           |
|                                                                           |
| Channel subsystem ID : 3                                                  |
| Channel path ID  . . : B1     Channel path type  . : OSC                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 3      MENSA36          6      OS    RDBKAIWK                           |
| / 3      MENSA37          7      OS    RDBKESI3                           |
| / 3      MENSA38          8      OS    WTSC01                             |
| / 3      MENSA39          9      OS    WTSCNET                            |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 13-2 Processors: Define Access List - OSC

In this example, because we select all partitions to the Access List, we do not see the Define Candidate List panel, even though we defined the CHPID as SHR.

HCD now returns to the Channel Path List, and you see the CHPID that was defined (see Figure 13-3).

```
-------------------------------------------------------------------------------
                              Channel Path List     Filter Mode. More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ B1    1D4    OSC   SHR   __  __ __       No  FC#0454 OSA-Express7S 1.2 GbE LX
******************************* Bottom of data ********************************
```

Figure 13-3 Processors: Channel Path List - OSC

13.2.2 Defining OSC CHPID connections to an OSC control unit

OSC Control Unit numbers come from the common pool of Hex '0000' to 'FFFE'. The HCD default unit address range for the 254 devices is Hex '00' to 'FD'. Devices with unit addresses FE and FF are not allowed.

The only way to define an Open Systems Adapter (OSA) connection to its CU is direct connect.

Here are considerations for connecting an OSC CHPID to an OSC CU and its 3270-X devices:

- Determine how many OSCs are required to provide a primary and secondary/backup network connection.
- The example connects to a predefined OSC CU (1B10) and 3270-X devices 0880 - 088F.

To define OSC CHPID connections to an OSC CU, complete the following steps:

1. From the main HCD panel, select option `1.4. Control units`.
2. Scroll through the CU list until you find the CU that you want to connect to, or in the command line, enter `L 1B10`. In our example, we use `1B10`.
3. Enter `c` next to the CU definition, and press Enter.
4. Changes the CU definition that you want (see Figure 13-4), and press Enter.

```
+----------------- Change Control Unit Definition ------------------+
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . 1B10  +                              |
|                                                                   |
| Control unit type  . . . . . OSC           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . ________________________________     |
|                                                                   |
| Connected to switches  . . . __  __  __  __  __  __  __  __  +    |
| Ports  . . . . . . . . . . . __  __  __  __  __  __  __  __  +    |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
+-------------------------------------------------------------------+
```

Figure 13-4 Control units: Change Control Unit Definition - OSC

5. Enter `c` next to the `Processor.CSS` that contains the partition that you want to have access to the CU and also has access to the CHPID you want to connect to the CU. Then, press Enter.
6. Update `Channel path IDs` to `B1` to define the Processor / CU connection (see Figure 13-5 on page 270). CHPID B1 is the CHPID that we defined in the previous example. Press Enter.

```
+--------------------- Change Control Unit Definition ----------------------+
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Control unit number  . : 1B10        Type . . . . . . : OSC               |
| Processor ID . . . . . : MENSA       Mensa                                |
| Channel Subsystem ID . : 3                                                |
|                                                                            |
| Channel path IDs . . . . B1    __   __   __   __   __   __   __    +     |
| Link address . . . . . . ____ ____ ____ ____ ____ ____ ____ ____   +     |
|                                                                            |
| Unit address . . . . . . 00    __   __   __   __   __   __   __    +     |
| Number of units  . . . . 254   ___  ___  ___  ___  ___  ___  ___         |
|                                                                            |
| Logical address  . . . . __  + (same as CUADD)                           |
|                                                                            |
| Protocol . . . . . . . . __  + (D, S or S4)                              |
| I/O concurrency level  . _   + (1, 2 or 3)                               |
|                                                                            |
|                                                                            |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 13-5 Control units: Change Control Unit Definition detail - OSC

7. HCD now shows the Select Processor / CU panel. This panel is where the connection is now made between the OSC CHPID (B1) and the CU (1B10). Press Enter again to return to the Control Unit List screen.

13.2.3 Defining 3270-X devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines what devices a z/OS system can access when it undergoes an IPL. Also, the partition that the z/OS system is restarted in also must be able to access the CHPIDs that connect to the CUs and devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups, which are defined in Eligible Device Tables (EDTs) within an OSCONFIG. OSA definitions usually do not use Esoterics.

Here are considerations for adding devices to an OSCONFIG:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access that device.
- In this example, we add the predefined OSC devices 0880 - 088F (3270-X).

To define 3270-X devices to an OSCONFIG, complete the following steps:

1. From the main HCD panel, select option `1.5. I/O Devices`.
2. Scroll through the I/O Device List until you find the device number that you want to add to the OSCONFIG, or in the command line, enter `L 0880`. In our example, we use `0880`.
3. Enter `c` next to one or more device numbers, and press Enter.
4. HCD displays the Change Device Definition panel where you can modify the CU that the devices are attached to (see Figure 13-6 on page 271). Press Enter.

```
+------------------------- Change Device Definition --------------------------+
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Device number  . . . . . . . : 0880  (0000 - FFFF)                           |
| Number of devices  . . . . . : 16                                            |
| Device type  . . . . . . . . : OSC                                           |
|                                                                              |
| Serial number  . . . . . . . . __________  +                                 |
| Description  . . . . . . . . . ________________________________              |
|                                                                              |
| Volume serial number . . . . . ______  + (for DASD)                          |
|                                                                              |
| PPRC usage . . . . . . . . . . _  + (for DASD)                               |
|                                                                              |
| Connected to CUs . 1B10  ____  ____  ____  ____  ____  ____  ____  +         |
|                                                                              |
| ENTER to continue.                                                           |
|                                                                              |
+------------------------------------------------------------------------------+
```

Figure 13-6 I/O Devices: Change Device Definition - OSC

Note: A maximum of 120 valid subchannels can be used on an OSC channel path even though IOCP allows you to define more than this number. The Support Element manages this maximum when sessions are configured under Advanced Facilities.

5. HCD now displays the Device / Processor Definition panel where you may modify some of the device parameters relating to Subchannel Set (SS), UA, and the Explicit Device Candidate List (see Figure 13-7).

```
+---------------------- Device / Processor Definition -----------------------+
|                                                                 Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Select processors to change device/processor definitions, then press       |
| Enter.                                                                     |
|                                                                            |
| Device number  . . : 0880      Number of devices  . : 16                   |
| Device type  . . . : 3270-X                                                |
|                                                                            |
|                                          Preferred  Device Candidate List  |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null   |
| _ MENSA.3     0    00   No        No      __         No             ___    |
| ***************************** Bottom of data ***************************** |
+----------------------------------------------------------------------------+
```

Figure 13-7 I/O Devices: Device / Processor Definition continued - OSC

6. In the HCD panel, we define devices to the OSCONFIG. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG that you want to add to the devices to, or in the command line, enter `L MENSA32`. In our example, we use `MENSA32`.
7. Enter `s` next to the OSCONFIG, and press Enter.

HCD displays the device parameters and features that are applicable to that device type. In our example, we add 3270-X devices to MENSA32.

8. Make the following updates to define the Device Parameter (see Figure 13-8), and press Enter:
   - Update OFFLINE to No (if you want these devices to be online during IPL time).
   - Update DYNAMIC to Yes (if you want the device to be dynamically changeable).
   - Update LOCANY to Yes (if the device UCB can be in 31-bit storage).
   - Update DOCHAR to Yes (if you want to use the US character set).

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                 Row 1 of 22 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise the values below.                                         |
|                                                                             |
| Configuration ID . : MENSA32      Z17 CONFIG SETUP                          |
| Device number  . . : 0880         Number of devices  : 16                   |
| Device type  . . . : 3270-X                                                 |
|                                                                             |
| Parameter/                                                                  |
| Feature    Value +          R Description                                   |
| OFFLINE    No                  Device considered online or offline at IPL   |
| DYNAMIC    Yes                 Device has been defined to be dynamic        |
| LOCANY     Yes                 UCB can reside in 31 bit storage             |
| ASCACHAR   No                  ASCII A Character Generator                  |
| ASCBCHAR   No                  ASCII B Character Generator                  |
| DOCHAR     Yes                 United States English Character Generator    |
| FRCHAR     No                  French Character Generator                   |
| GRCHAR     No                  German Character Generator                   |
| KACHAR     No                  Katakana Character Generator                 |
| UKCHAR     No                  United Kingdom English Character Generator   |
| AUDALRM    No                  Audible Alarm                                |
+-----------------------------------------------------------------------------+
```

Figure 13-8 I/O Devices: Define Device Parameters / Features - OSC

9. The Assign/Unassign Device to Esoteric panel opens. You can specify which Esoteric (if any) that you want the devices to be added to. Press Enter (see Figure 13-9 on page 273).

   In this example, we add only the OSC/3270-X devices to the OSCONFIG MENSA32 and not to any Esoterics.

```
+----------- Define Device to Operating System Configuration -----------+
|                                                            Row 1 of 11 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                        |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                        |
| Device number  . : 0880          Number of devices  : 16             |
| Device type  . . : 3270-X                                             |
|                                                                        |
| Subchannel Sets used in processor configurations  : 0                 |
|                                                                        |
| / Config. ID   Type     SS Description                    Defined     |
| _ ITSO         MVS      0  All ITSO devices               Yes         |
| _ ITSOBLD      MVS      0  ITSO + CPO devices            Yes         |
| _ MENSA32      MVS      0  Z17 CONFIG SETUP              Yes         |
| _ PRF1         MVS      0  Proof Points                  Yes         |
| _ WAT1         MVS      0  Watsonx                       Yes         |
| _ ZOSMAINT     MVS      0  z/OS Central Mgmt Image       Yes         |
| _ ZOSMESUR     MVS      0  z/OS systems to monitor all   Yes         |
| _ ZOS24BAS     MVS      0  z/OS 2.4 base system          Yes         |
| _ ZOS24P1K     MVS      0  z/OS 2.4 P1K study system     Yes         |
| _ ZOS25BAS     MVS      0  z/OS 2.5 inital system        Yes         |
| _ ZOS25POP     MVS      0  z/OS 2.5 POP system           Yes         |
+------------------------------------------------------------------------+
```

Figure 13-9 I/O Devices: Define Device to Operating System Configuration - OSC

10. The final panel opens and shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

13.2.4 Defining 3270-X devices to the Nucleus Initialization Program (NIP) Console List within an OSCONFIG

The NIP Console List determines the device addresses that are eligible to receive NIP or IPL messages in the early startup stages of when z/OS is started.

The devices must first be defined to an OSCONFIG so that they can be added to a NIP Console List within an OSCONFIG.

The NIP Console List also determines which console receives the NIP/IPL messages first. If that console is unavailable, then the NIP tries the next device in the list until all devices in the list are tried.

If the NIP cannot write IPL messages to any 3270-X device in the list, then the messages are written to the Hardware Management Console (HMC) Operating System Messages panel.

To view these messages, complete the following steps:

1. Select the logical partition (LPAR) for IPL on the HMC.
2. Click the **>>** breakout symbol next to the LPAR name.
3. Select **Daily** → **Operating System Messages**, as shown in see Figure 13-10 on page 274.

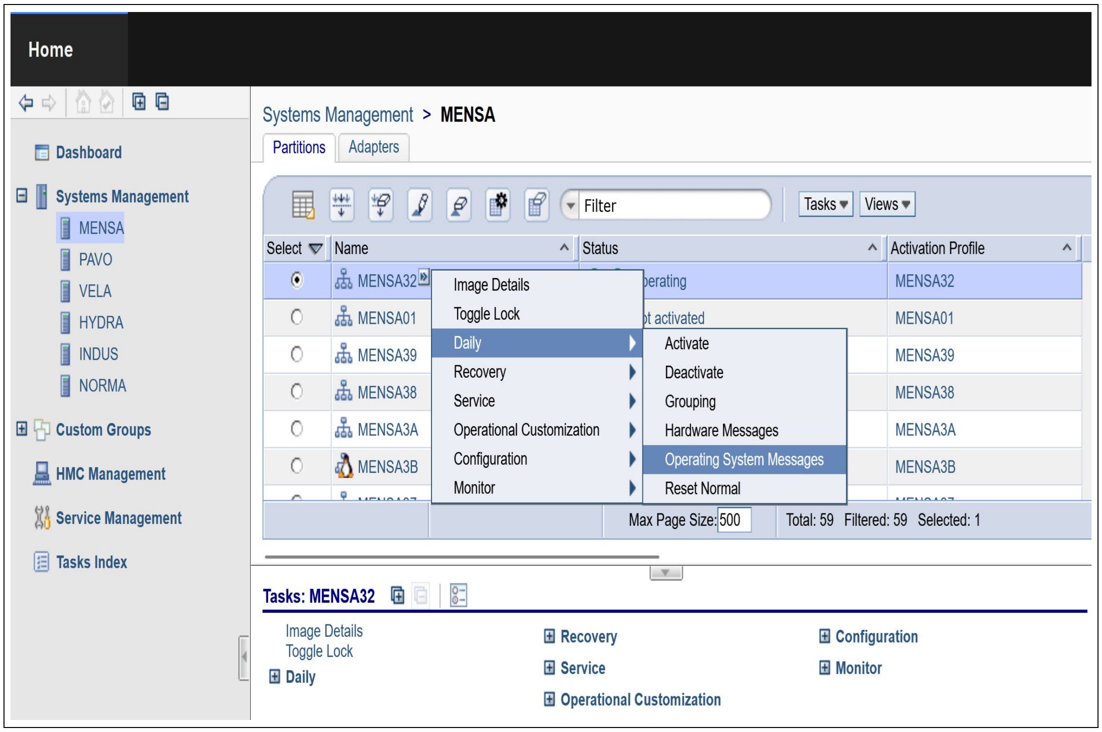

Figure 13-10 Operating System Messages

Commands and displays can be entered into the Command field (see Figure 13-11).

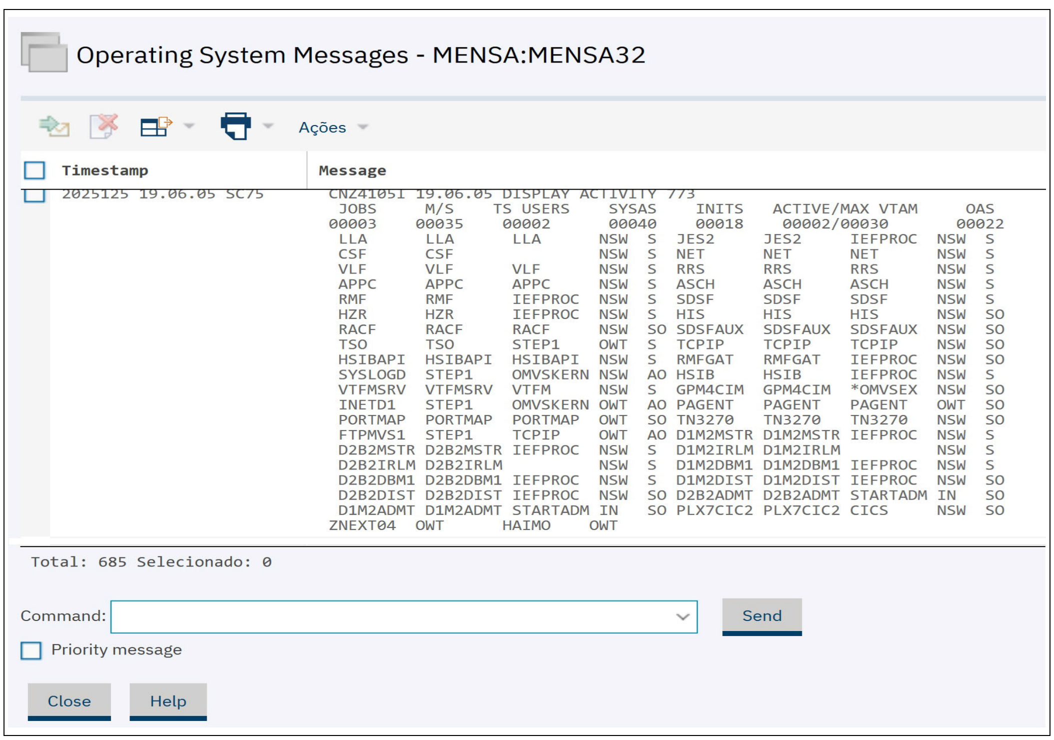

Figure 13-11 Operating System Messages command interface

Here are considerations for adding devices to a NIP Console List within an OSCONFIG:

- Adding a device to a NÌP Console List within an OSCONFIG does not necessarily mean that the NIP may write IPL messages to that device.
- The devices that are defined in the NIP Console List also need CU and CHPID access to the partition where z/OS is being started.
- On the HMC under OSA Advanced Facilities, the OSC (Open Systems Adapter Integrated Console Controller (ICC) (OSA-ICC) console Server and Session definitions must be defined and activated.
- A valid 3270-X session (that uses IBM® Personal Communications or an equivalent 3270 emulator) must also be connected to the OSA-ICC Session. This configuration enables a valid session to be established to the OSA-ICC for NIP messages to be delivered to that device.
- In this example, we add the predefined OSC devices 0880-088F (3270-X).

To define 3270-X devices to the NIP within an OSCONFIG, complete the following steps:

1. From the main HCD panel, select option `1.1. Operating system configurations` to display the Operating System Configuration List.
2. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG that you want to add to the 3270-X devices to the NIP Console List, or in the command line, enter `L MENSA32`. In our example, we use `MENSA32`.
3. Enter `n` next to the OSCONFIG, and press Enter.

   HCD displays the defined devices in the NIP Console List (see Figure 13-12).

```
+--------------------------- NIP Console List ----------------------------+
|   Goto  Backup  Query  Help                                             |
| ----------------------------------------------------------------------  |
|                                                                         |
| Command ===> _______________________________________ Scroll ===> CSR    |
|                                                                         |
| Select one or more consoles, then press Enter. To add, use F11.         |
|                                                                         |
| Configuration ID . : MENSA32      Z17 - CONFIG SETUP                    |
|                                                                         |
|   Order   Device                                                        |
| / Number  Number  Device Type                                           |
| *************************** Bottom of data **************************** |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
```

Figure 13-12 OCSONFIGs: NIP Console List

In our example, there are no devices that are defined in the NIP Console List.

4. In the command line, enter `add` (see Figure 13-13) to add a 3270-X device to the NIP Console List.
5. Update `Device number of console` to `0880`, and press Enter.

```
+--------------------------- NIP Console List ----------------------------+
|   Goto  Backup  Query  Help                                             |
| ----------------------------------------------------------------------  |
|                                                                         |
| Command ===> _______________________________________ Scroll ===> CSR    |
|                                                                         |
| Select one or more consoles, then press Enter. To add, use F11.         |
|                                                                         |
| Configuration ID . : MENSA32      Z17 CONFIG SETUP                      |
|                  +-------------- Add NIP Console ---------------+       |
|   Order   Device |                                              |       |
| / Number  Number |                                              |       |
| **************** | Specify the following values.                | ***** |
|                  |                                              |       |
|                  | Device number of console . . . . . . 880_    |       |
|                  |                                              |       |
|                  | Order number . . . . . . . . . . . . 1       |       |
|                  |                                              |       |
|                  |                                              |       |
|                  |                                              |       |
|                  +----------------------------------------------+       |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
```

Figure 13-13 OSCONFIGs: Add NIP Console

Because this device entry is the first one in the list, the order is 1 (see Figure 13-14).

```
+--------------------------- NIP Console List ----------------------------+
|   Goto  Backup  Query  Help                                             |
| ----------------------------------------------------------------------  |
|                                                            Row 1 of 1   |
| Command ===> _______________________________________ Scroll ===> CSR    |
|                                                                         |
| Select one or more consoles, then press Enter. To add, use F11.         |
|                                                                         |
| Configuration ID . : MENSA32      Z17 CONFIG SETUP                      |
|                                                                         |
|   Order   Device                                                        |
| / Number  Number  Device Type                                           |
| _   1     0880    3270-X                                                |
| *************************** Bottom of data **************************** |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
```

Figure 13-14 OCSONFIGs: NIP Console added

6. Add device 0881 to the NIP Console List (see Figure 13-15).

```
+--------------------------- NIP Console List ----------------------------+
|   Goto  Backup  Query  Help                                              |
| ----------------------------------------------------------------------   |
|                                                              Row 1 of 2 |
| Command ===> _______________________________________ Scroll ===> CSR     |
|                                                                          |
| Select one or more consoles, then press Enter. To add, use F11.          |
|                                                                          |
| Configuration ID . : MENSA32      Z17 CONFIG SETUP                       |
|                                                                          |
|   Order   Device                                                         |
| / Number  Number  Device Type                                            |
| _   1     0881    3270-X                                                 |
| _   2     0880    3270-X                                                 |
| ************************** Bottom of data *************************** |
+--------------------------------------------------------------------------+
```

Figure 13-15 OCSONFIGs: Extra NIP Console added

13.2.5 Verifying OSC definition

More details about how to verify OSC definition, see 7.4, “Verifying the OSA-ICC definition” on page 151.

13.3 OSD definitions (Queued Direct I/O - QDIO)

This section covers defining OSD CHPID, control unit and devices.

13.3.1 Defining OSD CHPIDs

Queued Direct I/O (QDIO) is a highly efficient data transfer architecture, which dramatically improves data transfer speed and efficiency for TCP/IP traffic. QDIO mode is referred to as OSD because the CHPID type coded in the IOCDS is OSD.

Note: The IBM z17 introduces a new adapter called Network Express that can be viewed as the next generation regarding OSA Express. IBM recommends migrating OSD channels to the new channel type OSH.

When defining an OSD connection, first you must determine which type of OSA-Express feature you need for your configuration (see Table 13-2):

Table 13-2 IBM z17 - New Build (NB) & Carry Forward (CF)

Networking Features	Feature Code (FC)	New Build (NB)/Carry Forward (CF)
OSA Express 7S GbE LX	0442	CF from z15 only
OSA Express 7S GbE SX	0443	CF from z15 only
OSA Express 7S 10GbE LR	0444	CF from z15 only
OSA Express 7S 10GbE SR	0445	CF from z15 only

Networking Features	Feature Code (FC)	New Build (NB)/Carry Forward (CF)
OSA Express7S 1.2 SX 1G	0455	CF-NB
OSA Express7S 1.2 LX 1G	0454	CF-NB
OSA Express7S 1.2 SR 10G	0457	CF-NB
OSA Express7S 1.2 LR 10G	0456	CF-NB
OSA Express7S 1.2 SR 25G	0459	CF-NB
OSA Express7S 1.2 LR 25G	0460	CF-NB
OSA Express7S 1000BaseT	0446	CF from z15 only

Note: Starting on z17, OSA Express7S 1000BaseT will support OSD channel type only.

Here are the considerations for a new OSD CHPID:

- Channel path id (CHPID).
- Channel ID (CHID).
- Channel path type.
- Operational mode.
- Description.
- Partition access list.
- For performance and redundancy, determine how many I/O cards of that feature are installed in the processor and to what PCIe ports on what CPC drawer do the I/O cards connect to. (For a list of installed hardware, see the PCHID/CHID report.)

To define an OSD CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. In the command line, enter `add` to add a CHPID. (see Figure 13-16 on page 279)
5. Make the following updates, and press Enter:
   - Update `Channel path ID` to `D8`.
   - Update `Channel ID` to `174`.
   - Update `Channel path type` to `OSD`.
   - Update `Operational mode` to `SHR`. (To share this channel path among logical partitions or even SPAN to share over multiple CSSs).
   - Update `Description` to the description that you want.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA       Mensa                                  |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . D8    +           Channel ID  174  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSD   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0445 OSA-Express7S 10 GbE SRX                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
+---------------------------------------------------------------------------+
```

Figure 13-16 Processors: Add Channel Path - OSD

6. HCD now prompts you to specify `Will greater than 160 TCP/IP stacks be required for this channel`? The default is `No`, which we use for our example (see Figure 13-17). For more information see the *IBM System z Connectivity Handbook*, SG24-5444. Press Enter.

```
+----------- Allow for more than 160 TCP/IP stacks -----------+
|                                                             |
|                                                             |
| Specify Yes to allow more than 160 TCP/IP stacks,           |
| otherwise specify No. Specifying Yes will cause priority    |
| queuing to be disabled.                                     |
|                                                             |
| Will greater than 160 TCP/IP stacks                         |
| be required for this channel?  . . . No                     |
|                                                             |
|                                                             |
+-------------------------------------------------------------+
```

Figure 13-17 Processors: Allow for more than 160 TCP/IP stacks - OSD

7. Next, HCD now prompts you to add or modify any physical network IDs. We do not use any physical network ID definitions for the OSD definition in this example.
8. Leave the default option for Physical Network IDs as blank fields (see Figure 13-18 on page 280), and press Enter.

```
*--------------------- Add/Modify Physical Network IDs ---------------------*
|                                                                          |
|                                                                          |
| If the Channel ID (CHID) is associated to one or more physical networks, |
| specify each physical network ID corresponding to each applicable        |
| physical port.                                                           |
|                                                                          |
| Physical network ID 1  . . ________________                              |
| Physical network ID 2  . . ________________                              |
| Physical network ID 3  . . ________________                              |
| Physical network ID 4  . . ________________                              |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
*--------------------------------------------------------------------------*
```

Figure 13-18 Processors: Add/Modify Physical Network IDs - OSD

9. HCD prompts you to select the partitions that will have access to the CHPID. Enter a forward slash (/) next to the partition that you want (see Figure 13-19), and press Enter.

```
+--------------------------- Define Access List ----------------------------+
|                                                            Row 12 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                          |
| Select one or more partitions for inclusion in the access list.          |
|                                                                          |
| Channel subsystem ID : 3                                                 |
| Channel path ID  . . : D8     Channel path type  . : OSD                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| / 3      MENSA36          6      OS    RDBKAIWK                          |
| / 3      MENSA37          7      OS    RDBKESI3                          |
| / 3      MENSA38          8      OS    WTSC01                            |
| / 3      MENSA39          9      OS    WTSCNET                           |
| **************************** Bottom of data ***************************** |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
+--------------------------------------------------------------------------+
```

Figure 13-19 Processors: Define Access List - OSD

Because we select all partitions to the access list, we do not see the Define Candidate List panel, even though we defined the CHPID as `SHR`.

The HCD now returns to the Channel Path List panel and shows you the CHPID that was defined (see Figure 13-20),

```
--------------------------------------------------------------------------------
                              Channel Path List     Filter Mode. More:      >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ D8    174   OSD   SHR   __  __ __       No  FC#0445 OSA-Express7S 10 GbE SRX
******************************* Bottom of data ********************************
```

Figure 13-20 Processors: Channel Path List - OSD

13.3.2 Defining OSD CHPID connections to an OSA control unit

All OSA QDIO channels support 480 unit addresses (subchannels). While the 480 subchannels can be distributed across many LPARs, only 255 unit addresses can be assigned to one LPAR. This is because only one control unit is defined on the OSA port. One control unit can only accommodate 256 devices per LPAR. (One device is reserved to hardware, the X'FF' unit address.)

The only way to define an OSA connection to its CU is direct connect.

You might want to connect the OSA CU definition to multiple CPCs even though the physical OSA is still unique to any one CPC. Also, you might want to span the OSA over multiple CSSs within a CPC.

Here are considerations for connecting an OSD CHPID to an OSA CU and its OSA devices:

- Determine how many OSAs are required to provide a primary and secondary/backup network connection.
- The example connects to a predefined OSA CU (1D80), OSA devices 1D80 - 1D8F.

To define OSD CHPID connections to an OSA CU, complete the following steps:

1. From the main HCD panel, select option `1.4. Control units`.
2. Scroll through the CU list until you find the CU that you want to connect to, or in the command line, enter `L 1D80`. In our example, we use `1D80`.
3. Enter `c` next to the CU definition, and press Enter.

4. Change the CU definition as you want (see Figure 13-21), and press Enter.

```
+----------------- Change Control Unit Definition ------------------+
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . 1D80  +                              |
|                                                                   |
| Control unit type  . . . . . OSA           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . FC#0445 OSA-Express7S 10 GbE SRX     |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
+-------------------------------------------------------------------+
```

Figure 13-21 Control units: Change Control Unit Definition - OSD

5. Enter `c` next to the `Processor.CSS` that contains the partition that you want to access the CU and also access the CHPID that you want to connect to the control unit. Press Enter.
6. Update `Channel path IDs` to `D8` to define the Processor / CU connection (see Figure 13-22). CHPID D8 is the CHPID that we defined in the previous example. Press Enter.

```
+---------------------- Change Control Unit Definition -----------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : 1D80        Type . . . . . . : OSA                 |
| Processor ID . . . . . : MENSA       Mensa                                  |
| Channel Subsystem ID . : 3                                                  |
|                                                                             |
| Channel path IDs . . . . D8    __    __    __    __    __    __    __    +  |
| Link address . . . . . . ____  ____  ____  ____  ____  ____  ____  ____  +  |
|                                                                             |
| Unit address . . . . . . 00    __    __    __    __    __    __    __    +  |
| Number of units  . . . . 255   ___   ___   ___   ___   ___   ___   ___      |
|                                                                             |
| Logical address  . . . . __  + (same as CUADD)                              |
|                                                                             |
| Protocol . . . . . . . . __  + (D, S or S4)                                 |
| I/O concurrency level  . _   + (1, 2 or 3)                                  |
|                                                                             |
+-----------------------------------------------------------------------------+
```

Figure 13-22 Control units: Change Control Unit Definition detail - OSD

> **Note:** For OSD configurations, unit address FF is not allowed.

7. HCD now shows the Select Processor / CU panel, where the connection is now made between the OSD CHPID (D8) and the CU (1D80). See Figure 13-23 and Press Enter. Press Enter again to return to the Control Unit List screen.

```
 Select Processor / CU    Row 1 of 15 More:       >
Command ===> ______________________________________________ Scroll ===> CSR

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

Control unit number . . : 1D80     Control unit type . . . : OSA

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ MENSA.3    D8      _______ _______ _______ _______ _______ _______ _______
_ MENSA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.1    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.4    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.5    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-23 Control units: `Select Processor / CU` *- OSD*

13.3.3 Defining OSA devices to an OSCONFIG(OSD CHIPDs)

The OSCONFIG name is the part of an IODF that determines what devices a z/OS system may access when it undergoes an IPL. Also, the partition that the z/OS system is started in also must have access to the CHPIDs that connect to the CUs and devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups, which are defined in EDTs within an OSCONFIG. OSA definitions usually do not use Esoterics.

Here are considerations for adding devices to an OSCONFIG:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system has access to that device.
- In this example, we add the predefined OSA devices 1D80 - 1D8F (OSA).

To define OSA devices to an OSCONFIG, complete the following steps:

1. From the main HCD panel, select option `1.5. I/O Devices`.
2. Scroll through the I/O Device List until you find the device number that you want to add to the OSCONFIG, or in the command line, enter `L 1D80`. In our example, we use `1D80`.
3. Enter `c` next to one or more device numbers, and press Enter.
4. HCD displays the Change Device Definition panel, where you may modify the CU that the devices are attached to (see Figure 13-24 on page 284). Press Enter.

```
+------------------------- Change Device Definition --------------------------+
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Device number  . . . . . . . : 1D80  (0000 - FFFF)                           |
| Number of devices  . . . . . : 16                                            |
| Device type  . . . . . . . . : OSA                                           |
|                                                                              |
| Serial number  . . . . . . . . __________  +                                 |
| Description  . . . . . . . . . ________________________________              |
|                                                                              |
| Volume serial number . . . . . ______  + (for DASD)                          |
|                                                                              |
| PPRC usage . . . . . . . . . . _  + (for DASD)                               |
|                                                                              |
| Connected to CUs . 1D80  ____  ____  ____  ____  ____  ____  ____  +         |
|                                                                              |
| ENTER to continue.                                                           |
|                                                                              |
|                                                                              |
|                                                                              |
+------------------------------------------------------------------------------+
```

Figure 13-24 I/O Devices: Change Device Definition - OSD

► **Note:** OSD definition without the OSAD device (Unit Address FE). OSA Advanced Facilities on the HMC does not need OSAD device. This device was used to communicate between the OSA and the operational system running OSA/SF.

5. HCD now displays the Device / Processor Definition panel, where you may modify some of the Device parameters relating to SS, UA, and the Explicit Device Candidate List. Type `c` next to the `Proc.CCSID` item if you want to modify any of these parameters (see Figure 13-25), or press Enter to accept the defaults.

```
+---------------------- Device / Processor Definition -----------------------+
|                                                               Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR   |
|                                                                            |
| Select processors to change device/processor definitions, then press      |
| Enter.                                                                     |
|                                                                            |
| Device number  . . : 1D80       Number of devices  . : 16                  |
| Device type  . . . : OSA                                                   |
|                                                                            |
|                                         Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null  |
| _ MENSA.3     0    00   No        No      __         No             ___   |
| **************************** Bottom of data **************************** |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 13-25 I/O Devices: Device / Processor Definition continued - OSD

6. The HCD panel opens, where we define devices to the OSCONFIG. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG that you want to add the devices to, or in the command line, enter `L MENSA32`. In our example, we use `MENSA32`.
7. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the device parameters and features that are applicable to that device type. In our example, we add OSA devices to `MENSA32`.
8. Make the following updates to define the device parameter (see Figure 13-26), and press Enter:
   – Update `OFFLINE` to `No` (if you want these devices to be online during IPL time).
   – Update `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   – Update `LOCANY` to `Yes` (if the device UCB can be in 31-bit storage).

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                   Row 1 of 3 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise the values below.                                         |
|                                                                             |
| Configuration ID . : MENSA32      Z17 CONFIG SETUP                          |
| Device number  . . : 1D80         Number of devices  : 16                   |
| Device type  . . . : OSA                                                    |
|                                                                             |
| Parameter/                                                                  |
| Feature    Value +          R Description                                   |
| OFFLINE    No                 Device considered online or offline at IPL    |
| DYNAMIC    Yes                Device has been defined to be dynamic         |
| LOCANY     Yes                UCB can reside in 31 bit storage              |
| **************************** Bottom of data ***************************** |
+-----------------------------------------------------------------------------+
```

Figure 13-26 I/O Devices: Define Device Parameters / Features - OSD

9. The Assign/Unassign Device to Esoteric panel opens, where you can specify which Esoteric (if any) that you want the devices to be added to. We add only the OSA devices to the OSCONFIG `MENSA32` and not to any Esoterics in this example. Press Enter (see Figure 13-27).

```
+----------- Define Device to Operating System Configuration -----------+
|                                                            Row 1 of 11 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : 1D80          Number of devices  : 16              |
| Device type  . . : OSA                                                |
|                                                                       |
| Subchannel Sets used in processor configurations  : 0                 |
|                                                                       |
| / Config. ID   Type     SS Description                  Defined       |
| _ MENSA32      MVS      0  Z17 CONFIG SETUP             Yes           |
| _ PRF1         MVS         Proof Points                               |
| _ WAT1         MVS         Watsonx                                    |
+-----------------------------------------------------------------------+
```

Figure 13-27 I/O Devices: Define Device to Operating System Configuration - OSD

10. The final panel opens and shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

13.3.4 Verifying the OSD definition

This section presents some commands that you can use to verify the OSD configuration. Before starting, the system programmer must activate the new I/O to make changes to a running configuration.

You can check your definition by using the following z/OS commands:

- **DISPLAY M=CHP(xx)**

  Check whether the `CHPID DESC` is displayed as `OSA DIRECT EXPRESS` (Figure 13-28).

```
D M=CHP(D8)
IEE174I 16.13.28 DISPLAY M 545
CHPID D8:  TYPE=11, DESC=OSA DIRECT EXPRESS, ONLINE
DEVICE STATUS FOR CHANNEL PATH D8
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
01D8 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 0174
OPERATING SPEED = 10G
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 13-28 OSD D M=CHP

- **DISPLAY M=DEV(xxxx)**

  Using this command, you can confirm the channel path to a device. The node descriptor information that is returned includes the emulated control unit (CU) 1D80. Also included is the 9175 machine type and serial number, as shown in Figure 13-29.

```
D M=DEV(1D80)
IEE174I 16.15.57 DISPLAY M 547
DEVICE 01D80   STATUS=ONLINE
CHP                   D8
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             1D80
INTERFACE ID          7400
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 001730.009.IBM.02.9175000B9FB8.D800
SCP TOKEN NED     = 001730.009.IBM.02.9175000B9FB8.D800
SCP DEVICE NED    = 001732.001.IBM.02.9175000B9FB8.D800
```

Figure 13-29 OSD D M=DEV(xxx)

- **DISPLAY U**

  Using this command, you can confirm the device number and the status. Check whether the device number and type are defined correctly (Figure 13-30). OSD uses 3 devices (Read, Write and Datapath).

```
D U,,,1D80,3
IEE457I 16.35.54 UNIT STATUS 585
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
1D80 OSA  A-BSY                                   0
1D81 OSA  A                                       0
1D82 OSA  A-BSY                                   0
```

Figure 13-30 OSD D U,,,device

13.4 Network Express for channel type OSH

The IBM z17 introduces a new adapter called Network Express that can be viewed as the next generation regarding OSA Express and RoCE adapters. In a single adapter, Network Express provides a new channel type OSH and the new PCI function NETH for SMC-R. Both NETH PCI function and an OSH CHPID can coexist on the same PCHID. This chapter will cover HCD definitions to OSH channel type only.

OSH exploits Enhanced QDIO (EQDIO) which offers better performance than QDIO reducing latency, CPU utilization, improving virtualization scale and simplifying all definitions. In addition, the new channel type OSH provides:

- All functions available with OSD.
- Has a short list of additional parameters reducing configuration and setup time.
- All static VTAM TRLE and many TCP/IP interface statements have been eliminated and no port parameter definition is required.
- Each interface is defined with a single device. Control read/write devices have been replaced with control queues.
- All OSH ports are optical and the default MTU is 9000 (Jumbo Frame).
- An OSH channel path is associated with a single port on a Network Express card. Each port has a unique PCHID.

More information about HCD definitions to PCI function NETH see 15.2.3, “Defining a NETH PCIe function” on page 339.

13.4.1 Defining OSH CHIPDs

When defining an OSH CHIPD, first you must determine which type of Network Express feature that you need for your configuration (See Table 13-3):

Table 13-3 IBM z17 Network Express Adapters

Feature	Feature	CHPID Type	OS Support
Network Express 10G LR	0525	OSH	z/OS, z/VM, z/TPF
Network Express 10G SR	0524	OSH	z/OS, z/VM, z/TPF
Network Express 25G LR	0527	OSH	z/OS, z/VM, z/TPF
Network Express 25G SR	0526	OSH	z/OS, z/VM, z/TPF

Here are the considerations for a new OSH CHPID:

- Channel path id (CHPID).
- Channel ID (CHID).
- Channel path type.
- Operational mode.
- Description.
- Partition access list.
- For performance and redundancy, determine how many I/O cards of that feature are installed in the processor and to what PCIe ports on what CPC drawer do the I/O cards connect to. (For a list of installed hardware, see the PCHID/CHID report.)

To define an OSH CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. In the command line, enter `add` (see Figure 13-31) to add a CHPID.
5. Make the following updates and press Enter:
   - Update `Channel path ID` to `91`.
   - Update `Channel ID` to `13C`.
   - Update `Channel path type` to `OSH`.
   - Update `Operational mode` to `SHR`. (To share this channel path among logical partitions or even SPAN to share over multiple CSSs).
   - Update `Description` to the description that you want.

```
+--------------------------- Add Channel Path ----------------------------+
|                                                                         |
|                                                                         |
| Specify or revise the following values.                                 |
|                                                                         |
| Processor ID . . . . : MENSA       Mensa                                |
| Configuration mode . : LPAR                                             |
| Channel Subsystem ID : 3                                                |
|                                                                         |
| Channel path ID  . . . . 91    +           Channel ID  13C  +           |
| Number of CHPIDs . . . . 1                                              |
| Channel path type  . . . OSH   +                                        |
| Operation mode . . . . . SHR   +                                        |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +    |
| Description  . . . . . . FC#0524 Network Express 10Gb SR_               |
|                                                                         |
| Specify the following values only if connected to a switch:             |
| Dynamic entry switch ID  __  + (00 - FF)                                |
| Entry switch ID  . . . . __  +                                          |
| Entry port . . . . . . . __  +                                          |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
```

Figure 13-31 Processors: Add Channel Path - OSH

6. HCD now prompts you to specify *Is the channel path used for link aggregation?* The default is No, which we use for our example (see Figure 13-32), and Press Enter.

```
+---------- Channel path used for link aggregation -----------+
|                                                              |
|                                                              |
| Specify Yes if the channel path is used for link             |
| aggregation, otherwise specify No.                           |
|                                                              |
| Is the channel path used for link                            |
| aggregation? . . . . . . . . . . . . No                      |
|                                                              |
|                                                              |
|                                                              |
+--------------------------------------------------------------+
```

Figure 13-32 Processors: Channel path used for link aggregation - OSH

Note: Link aggregation is a z/VM feature that can combine multiple OSA networking ports into one logical connection to increase the total throughput beyond what a single link could achieve. It also can provide redundancy where all but one of the physical links in a link aggregation group can fail without losing connectivity.

Note: To exploit this functionality, a NETH function cannot co-exist with the OSH channel path on its PCHID.

7. Next, HCD now prompts you to add or modify any physical network IDs. We do not use any physical network ID definitions for the OSH definition in this example.
8. Leave the default option for Physical Network IDs as blank fields, (see Figure 13-33), and press Enter.

```
*--------------------- Add/Modify Physical Network IDs ---------------------*
|                                                                            |
|                                                                            |
| If the Channel ID (CHID) is associated to one or more physical networks,  |
| specify each physical network ID corresponding to each applicable         |
| physical port.                                                             |
|                                                                            |
| Physical network ID 1  . . ________________                                |
| Physical network ID 2  . . ________________                                |
| Physical network ID 3  . . ________________                                |
| Physical network ID 4  . . ________________                                |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 13-33 Processors: Add/Modify Physical Network IDs - OSH

Note: When both an OSH channel path and a NETH function are defined on the same port of a network adapter, z/OS Communications Server requires the PNETID keyword be specified and they must have the same pnetid. This example will not exploit NETH.

9. HCD prompts you to select the partitions that will have access to the CHPID. Enter a forward slash (/) next to the partition that you want (see Figure 13-34), and press Enter.

```
+-------------------------- Define Access List ---------------------------+
|                                                            Row 13 of 15 |
| Command ===> ________________________________________ Scroll ===> CSR   |
|                                                                         |
| Select one or more partitions for inclusion in the access list.         |
|                                                                         |
| Channel subsystem ID : 3                                                |
| Channel path ID  . . : 91     Channel path type  . : OSH                |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                  |
|                                                                         |
| / CSS ID Partition Name   Number Usage Description                      |
| / 3      MENSA37          7      OS    RDBKESI3                         |
| / 3      MENSA38          8      OS    WTSC01                           |
| / 3      MENSA39          9      OS    WTSCNET                          |
| **************************** Bottom of data **************************** |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
```

Figure 13-34 Processors: Define Access List - OSH

10. Because we select all partitions to the access list, we do not see the Define Candidate List panel, even though we defined the CHPID as SHR.

11. The HCD now returns to the Channel Path List panel and shows you the CHPID that was defined (see Figure 13-35).

```
-------------------------------------------------------------------------------
                              Channel Path List     Filter Mode. More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 91    13C   OSH   SHR   __  __ __       No  FC#0524 Network Express 10Gb SR
****************************** Bottom of data *******************************
```

Figure 13-35 Processors: Channel Path List - OSH

13.4.2 Defining OSH CHPID connections to an OSA control unit

A valid subchannel for an OSH device can have any unit address in the range 00-FE. the X'FF' unit address is reserved to hardware. maximum of 256 valid subchannels can be defined on an OSH channel path.

OSH connections to an OSA control unit is defined just like OSD. More details see 13.3.2, “Defining OSD CHPID connections to an OSA control unit” on page 281.

13.4.3 Defining OSA devices to an OSCONFIG(OSH CHIPDs)

Channel types OSH and OSD use the same device type OSA, but OSH requires only one device to be defined for a TCP/IP stack. OSD requires three devices. More details see 13.3.3, “Defining OSA devices to an OSCONFIG(OSD CHIPDs)” on page 283

> **Note:** OSAD functionality is not available with OSH channels.

13.4.4 Verifying the OSH definition

This section presents some commands that you can use to verify the OSH configuration. Before starting, the system programmer must activate the new I/O to make changes to a running configuration.

You can check your definition by using the following z/OS commands:

- **`DISPLAY M=CHP(xx)`**

  Check whether the `CHPID DESC` is displayed as `OSA HYBRID` (Figure 13-36).

```
D M=CHP(91)
IEE174I 16.44.59 DISPLAY M 601
CHPID 91:  TYPE=35, DESC=OSA HYBRID, ONLINE
DEVICE STATUS FOR CHANNEL PATH 91
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
0191 +  +  +  +  +  +  +  +  .  .  .  .  .  .  .  +
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 013C
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 13-36 OSH D M=CHP

- **`DISPLAY M=DEV(xxxx)`**

  Using this command, you can confirm the channel path to a device. The node descriptor information that is returned includes the emulated control unit (CU) 1910. Also included is the 9175 machine type and serial number, as shown in Figure 13-37 on page 292.

```
D M=DEV(1910)
IEE174I 17.42.39 DISPLAY M 646
DEVICE 01910   STATUS=ONLINE
CHP                   91
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             1910
INTERFACE ID          0004
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 710003.010.IBM.02.9175000B9F  .9100
SCP TOKEN NED     = 001730.010.IBM.02.9175000B9F  .9100
SCP DEVICE NED    = 001732.010.IBM.02.9175000B9F  .9100
```

Figure 13-37 OSH D M=DEV(xxx)

- **DISPLAY U**

  Using this command, you can confirm the device number and the status. Check whether the device number and type are defined correctly (Figure 13-38). OSH uses only one device.

```
D U,,,1910,3
IEE457I 16.53.04 UNIT STATUS 611
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
1910 OSA  A-BSY                                   0
1911 OSA  O                                       0
1912 OSA  O                                       0
```

Figure 13-38 OSH D U,,,device

13.5 IQD CHPIDs for HiperSockets

The HiperSockets function, also known as internal queued direct input/output (iQDIO or internal QDIO), is an integrated function of the firmware of the System z central processor complexes (CPCs). It provides users with attachment to high-speed logical local area networks (LANs) with minimal system and network resource usage. HiperSockets provides internal virtual local area networks, which are Internet Protocol networks in the System z system.

This section describes the process to define IQD CHPIDs and their CUs and devices.

13.5.1 Defining IQD CHPIDs

When you define HiperSockets, use the CHPID channel type IQD. IQD CHPID also requires an static virtual channel ID (VCHID) statement.

> **Note:** Starting on z17, the VCHID range has expanded and now is 0x400-0x7FF. Static VCHIDs (used for IQD and ISM) will be in the range 0x7C0 to 0x7FF. More information about ISM see 15.2.2, “Defining an ISM PCIe function” on page 334.

IQD CHPIDs are defined logically and internally to the processor and require no installed hardware (PCHIDs). However, a maximum of up to 32 high-speed virtual local area network (VLAN) attachments may be defined.

Each of the logical IQD VCHIDs can support only one CHPID, but the CHPIDs may be spanned across multiple CSSs.

Here are considerations for a new IQD CHPID:

- Channel path id (CHPID)
- Virtual Channel id (VCHID)
- Channel path type
- Operational mode
- Description
- Partition access list

To define an IQD CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Type `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Type `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. In the command line, enter `add` (see Figure 13-39 on page 294) to add a CHPID.
5. Make the following updates and press Enter:
   - Update `Channel path ID` to `F0`.
   - Update `Channel ID` to `7C1`.
   - Update `Channel path type` to `IQD`.
   - Update `Operational mode` to `SHR`. (To share this channel path among logical partitions or even SPAN to share over multiple CSSs).
   - Update `Description` to the description that you want.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID . . . . : MENSA        Mensa                                  |
| Configuration mode . : LPAR                                                |
| Channel Subsystem ID : 3                                                   |
|                                                                            |
| Channel path ID  . . . . F0    +            Channel ID  7C1  +             |
| Number of CHPIDs . . . . 1                                                 |
| Channel path type  . . . IQD   +                                           |
| Operation mode . . . . . SHR   +                                           |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +       |
| Description  . . . . . . IQD Internal Queued Direct Comms                  |
|                                                                            |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                   |
| Entry switch ID  . . . . __  +                                             |
| Entry port . . . . . . . __  +                                             |
|                                                                            |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 13-39 Processors: Add Channel Path - IQD

HCD now prompts you to specify the IQD Channel Parameters, where you set the maximum frame size in KB and what IQD function is used.

6. Leave the default option of `16` for the Maximum frame size, and select option `1. Basic HiperSockets` for the IQD function (see Figure 13-40). Press Enter.

```
+------------------- Specify IQD Channel Parameters --------------------+
|                                                                        |
|                                                                        |
| Specify or revise the values below.                                    |
|                                                                        |
| Maximum frame size in KB . . . . . . 16  +                             |
|                                                                        |
| IQD function . . . . . . . . . . . . 1   1.  Basic HiperSockets        |
|                                          2.  IEDN Access (IQDX)        |
|                                          3.  External Bridge           |
|                                                                        |
| Physical network ID  . . . . . . . . ________________                  |
|                                                                        |
|                                                                        |
|                                                                        |
|                                                                        |
+------------------------------------------------------------------------+
```

Figure 13-40 Processors: Specify IQD Channel Parameters - IQD

Note: Each CHPID has a configurable frame size (16 KB, 24 KB, 40 KB, or 64 KB) that supports optimizing per HiperSocket LAN for small packets versus large packets.

7. HCD prompts you to select the partitions that will have access to the CHPID. Type forward slash (/) next to the needed partitions (see Figure 13-41), and press Enter.

```
+-------------------------- Define Access List ---------------------------+
|                                                            Row 13 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR  |
|                                                                          |
| Select one or more partitions for inclusion in the access list.          |
|                                                                          |
| Channel subsystem ID : 3                                                 |
| Channel path ID  . . : F0     Channel path type  . : IQD                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| / 3      MENSA37          7      OS    RDBKESI3                          |
| / 3      MENSA38          8      OS    WTSC01                            |
| / 3      MENSA39          9      OS    WTSCNET                           |
| *************************** Bottom of data ***************************** |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
+--------------------------------------------------------------------------+
```

Figure 13-41 Processors: Define Access List - IQD

HCD now displays the Define Candidate List panel. In our example, we do not select any Candidate LPARs for this IQD CHPID. Press Enter.

The HCD returns to the Channel Path List and show you the CHPID that was defined (see Figure 13-42).

```
-----------------------------------------------------------------------------
                              Channel Path List     Filter Mode. More:      >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ F0    7C1   IQD   SHR   __  __ __       No  IQD Internal Queued Direct Comms
****************************** Bottom of data *******************************
```

Figure 13-42 Processors: Channel Path List - IQD

13.5.2 Defining IQD CHPID connections to an IQD control unit

The only way to define an IQD connection to its CU is direct connect. You might want to span the IQD CHPID over multiple CSSs within a CPC.

Here are considerations for connecting an IQD CHPID to an IQD CU and its IQD devices:

- Determine how many IQDs are required to provide the required HiperSocket bandwidth.
- In this example, we connect to a predefined IQD CU (F001) and IQD devices F000 - 701F.

To define IQD CHPID connections to an IQD CU, complete the following steps:

1. From the main HCD panel, select option `1.4. Control units`.
2. Scroll through the CU list until you find the CU that you want to connect to, or in the command line, enter `L F001`. In our example, we use `F001`.
3. Enter `c` next to the CU definition, and press Enter.
4. Change the CU definition that you want as needed (see Figure 13-43), and press Enter.

```
+----------------- Change Control Unit Definition ------------------+
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . F001  +                              |
|                                                                   |
| Control unit type  . . . . . IQD           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . .                                      |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
+-------------------------------------------------------------------+
```

Figure 13-43 Control units: Change Control Unit Definition - IQD

5. HCD now shows the Select Processor / CU panel. This panel is where the connection is made between the IQD CHPID (F0) and the CU (F001). Enter `c` next to the `Processor.CSS` that contains the partition that you want to access the CU and the CHPID that you want to connect to the CU. Press Enter.

6. Update CHPID F0 to define the Processor / CU connection (see Figure 13-44), and press Enter.

```
+--------------------- Change Control Unit Definition ----------------------+
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Control unit number  . : F001        Type . . . . . . : IQD                |
| Processor ID . . . . . : MENSA       Mensa                                 |
| Channel Subsystem ID . : 3                                                 |
|                                                                            |
| Channel path IDs . . . . F0   __   __   __   __   __   __   __   +         |
| Link address . . . . . . ____ ____ ____ ____ ____ ____ ____ ____ +         |
|                                                                            |
| Unit address . . . . . . 00   __   __   __   __   __   __   __   +         |
| Number of units  . . . . 256  ___  ___  ___  ___  ___  ___  ___            |
|                                                                            |
| Logical address  . . . . __  + (same as CUADD)                             |
|                                                                            |
| Protocol . . . . . . . . __  + (D, S or S4)                                |
| I/O concurrency level  . _   + (1, 2 or 3)                                 |
|                                                                            |
|                                                                            |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 13-44 Control units: Change Control Unit Definition detail - IQD

Note: The control unit address range defines the number of devices that you can define for this control unit. You can define a smaller address range, but 256 gives you the maximum number of devices per control unit. If you need to define more than 256 devices per channel, you need to define more control units for that channel

7. HCD returns to the Select Processor / CU panel and shows the CHPID (F0) connection definition (see Figure 13-45). Press Enter again to return to the Control Unit List screen.

```
 . . . . . . . . . . . . . . . . . . . . . . . . . . .
                            Select Processor / CU    Row 1 of 15 More:      >
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : F001     Control unit type . . . : IQD

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ MENSA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.1    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MENSA.3    F0      _______ _______ _______ _______ _______ _______ _______
```

Figure 13-45 Control units: Select Processor / CU - IQD

13.5.3 Defining IQD devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines what devices a z/OS system has access to when it undergoes an IPL. Also the partition that the z/OS system is started in must have access to the CHPIDs that connect to the CUs and devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups that are defined in EDTs within an OSCONFIG. IQD definitions usually do not use Esoterics.

Here are considerations for adding devices to an OSCONFIG:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system has access to that device.
- You can define up to 256 devices per control unit.
- In this example, we add the predefined IQD devices F000 - F01F (IQD).

To define IQD devices to an OSCONFIG, complete the following steps:

1. From the main HCD panel, select option `1.5. I/O Devices`.
2. Scroll through the I/O Device List until you find the device number that you want to add to the OSCONFIG, or in the command line, enter `L F000`. In our example, we use `F000`.
3. Enter `c` next to one or more device numbers, and press Enter.
4. HCD displays the Change Device Definition panel, where you may modify the control unit that the devices are attached to (see Figure 13-46). Press Enter.

```
+------------------------ Change Device Definition -------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Device number  . . . . . . . : F000  (0000 - FFFF)                        |
| Number of devices  . . . . . : 32                                         |
| Device type  . . . . . . . . : IQD                                        |
|                                                                           |
| Serial number  . . . . . . . __________  +                                |
| Description  . . . . . . . . ________________________________             |
|                                                                           |
| Volume serial number . . . . . ______  + (for DASD)                       |
|                                                                           |
| PPRC usage . . . . . . . . . . _  + (for DASD)                            |
|                                                                           |
| Connected to CUs . F001  ____  ____  ____  ____  ____  ____  ____  +      |
|                                                                           |
| ENTER to continue.                                                        |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 13-46 I/O Devices: Change Device Definition - IQD

5. HCD now displays the Device / Processor Definition panel, where you may modify some of the device parameters relating to SS, UA, and the Explicit Device Candidate List. Enter `c` next to the `Proc.CCSID` item if you want to modify any of these parameters (see Figure 13-47), or press Enter to accept the defaults.

```
+---------------------- Device / Processor Definition -----------------------+
|                                                                 Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Select processors to change device/processor definitions, then press       |
| Enter.                                                                     |
|                                                                            |
| Device number  . . : F000       Number of devices  . : 32                  |
| Device type  . . . : IQD                                                   |
|                                                                            |
|                                          Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit        Null  |
| _ MENSA.3     0    00   No        No      __         No              ___   |
| **************************** Bottom of data ****************************  |
|                                                                            |
+----------------------------------------------------------------------------+
```

Figure 13-47 I/O Devices: Device / Processor Definition continued - IQD

6. The HCD panel where we define devices to the OSCONFIG is displayed. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG that you want to add to the devices to, or in the command line, enter `L MENSA32`. In our example, we use `MENSA32`.
7. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the device parameters and features that are applicable to that device type. In our example, we add IQD devices to `MENSA32`.
8. Make the following updates to define the Device Parameter (see Figure 13-48), and press Enter:
   – Update `OFFLINE` to `No` (if you want these devices to be online during IPL time).
   – Update `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   – Update `LOCANY` to `Yes` (if the device UCB can be in 31-bit storage).

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                  Row 1 of 3 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise the values below.                                         |
|                                                                             |
| Configuration ID . : MENSA32      Z17 - CONFIG SETUP                        |
| Device number  . . : F000         Number of devices  : 32                   |
| Device type  . . . : IQD                                                    |
|                                                                             |
| Parameter/                                                                  |
| Feature    Value +         R Description                                    |
| OFFLINE    No                Device considered online or offline at IPL     |
| DYNAMIC    Yes               Device has been defined to be dynamic          |
| LOCANY     Yes               UCB can reside in 31 bit storage               |
| **************************** Bottom of data ***************************** |
+-----------------------------------------------------------------------------+
```

Figure 13-48 I/O Devices: Define Device Parameters / Features - IQD

9. The Assign/Unassign Device to Esoteric panel is displayed, where you can specify which Esoteric (if any) that you want the devices to be added to. We add only the IQD devices to the OSCONFIG `MENSA32` and not to any Esoterics in this example. Press Enter (see Figure 13-49).

```
+----------- Define Device to Operating System Configuration -----------+
|                                                           Row 1 of 11 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : F000          Number of devices  : 32             |
| Device type  . . : IQD                                               |
|                                                                       |
| Subchannel Sets used in processor configurations  : 0                |
|                                                                       |
| / Config. ID   Type     SS Description                    Defined    |
| _ ITSO         MVS         All ITSO devices                           |
| _ ITSOBLD      MVS         ITSO + CPO devices                         |
| _ MENSA32      MVS      0  Z17 - CONFIG SETUP                Yes     |
| _ PRF1         MVS         Proof Points                               |
| _ WAT1         MVS         Watsonx                                    |
+-----------------------------------------------------------------------+
```

Figure 13-49 I/O Devices: Define Device to Operating System Configuration - IQD

10. The final panel opens and shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

13.5.4 Verifying the IQD definition

This section presents some commands that you can use to verify the IQD configuration. Before starting, the system programmer must activate the new I/O to make changes to a running configuration.

You can check your definition by using the following z/OS commands:

- **`DISPLAY M=CHP(xx)`**

  Check whether the `CHPID DESC` is displayed as `INTERNAL QUEUED DIRECT COMM` (Figure 13-50).

```
D M=CHP(F0)
IEE174I 18.20.29 DISPLAY M 740
CHPID F0:  TYPE=24, DESC=INTERNAL QUEUED DIRECT COMM, ONLINE
DEVICE STATUS FOR CHANNEL PATH F0
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
0F00 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
0F01 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
SWITCH DEVICE NUMBER = NONE
ATTRIBUTES = MFS(16KB)
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 13-50 IQD D M=CHP

- **`DISPLAY M=DEV(xxxx)`**

  Using this command, you can confirm the channel path to a device. The node descriptor information that is returned includes the emulated control unit (CU) F001 as shown in Figure 13-51.

```
D M=DEV(F000)
IEE174I 18.30.41 DISPLAY M 745
DEVICE 0F000   STATUS=ONLINE
CHP                   F0
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             F001
INTERFACE ID          F000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         =       .   .   .  .            .F000
SCP TOKEN NED     =       .   .   .  .            .F000
```

Figure 13-51 IQD D M=DEV(xxx)

- **`DISPLAY U`**

  Using this command, you can confirm the device number and the status. Check whether the device number and type are defined correctly (Figure 13-52). IQD uses 3 devices (Read, Write and Datapath).

```
D U,,,F000,3
IEE457I 18.33.03 UNIT STATUS 747
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
F000 OSA  A-BSY                                   0
F001 OSA  A                                       0
F002 OSA  A-BSY                                   0
```

Figure 13-52 IQD D U,,,device

14

Adding coupling connectivity

This chapter describes the steps to define CS5, CL5, CL6, and ICP channel path ID (CHPID) types, and the Coupling Facility (CF) links between these CHPID types.

It provides a list of these potential configuration items and a short description about how to do each of them by using a Hardware Configuration Definition (HCD).

Naming: The IBM z17 systems that are targeted by this publication consist of IBM z17 ME1. Throughout this chapter, we might refer to these machines as IBM z17.

Note: The examples that are shown in this chapter are based on the IBM z17 ME1 (9175).

Note: Not all the following configuration items are necessarily required for your installation. In addition, the examples that are presented are not exhaustive.

This chapter includes the following topics:

- Defining more I/O by using an HCD
- Coupling Facility logical partitions, CS5, CL5, CL6, and ICP CHPIDs

14.1 Defining more I/O by using an HCD

When defining new I/O components in an input/output definition file (IODF), certain definitions like operating system configurations (OSCONFIGs), partitions, Fibre Connection (FICON) switches, control units (CUs), and devices must be done first. After these items are defined, then the connections can be made:

- Defining Coupling Facility LPARs in a channel subsystem
- Defining CS5 CHPIDs
- Defining a Coupling Facility link with CS5 CHPIDs
- Defining CL5 and CL6 CHPIDs
- Defining a Coupling Facility link with CL5 or CL6 CHPIDs
- Defining ICP CHPIDs
- Defining a Coupling Facility link with ICP CHPIDs

The following I/O definitions use HCD to demonstrate the examples. The examples in this chapter continue the work example that was created in the previous chapters (for example, `SYS9.IODF81.WORK`).

14.2 Coupling Facility logical partitions, CS5, CL5, CL6, and ICP CHPIDs

This section covers defining CF logical partitions (LPARs) and the definitions for CS5, CL5, CL6, and ICP CHPID types.

14.2.1 Defining Coupling Facility LPARs in a channel subsystem

Here are considerations for a new (unreserved) partition:

- Partition name.
- Number.
- Usage.
- Description.
- To add CHPIDs to a partition, they first must be defined to the processor.
- Renaming an existing partition is a two-step process:
  a. Redefine the partition as reserved (`Partition name = *`), and activate the IODF on the processor.
  b. Redefine the partition with the new name, and activate the IODF on the processor.

To change a reserved partition to an active partition in a channel subsystem (CSS), complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a partition to, and press Enter.
3. Enter `p` next to the CSS ID that you want to add a partition to, and press Enter.
4. Enter `c` next to the Reserved Partition that you want to unreserve (we use `ID E` for this example), and press Enter.

5. Make the following updates (see Figure 14-1), and press Enter:
   - Update `Partition Name` to `MENSA3E` (a naming standard based on `CSS=3`, `Partition =E`).
   - Review `Partition usage` and change it if required. We use `CF` in this example.
   - Update `Description` to `Redbooks CF76`.

```
 *----------------------------- Partition List -----------------------------*
|   Goto  Backup  Query  Help                                              |
| ----------------------------------------------------------------------   |
|                                                                          |
| Command ===> ________________________________________ Scroll ===> CSR    |
|                                                                          |
| Select one or more partitions, then press Enter. To add, use F11.        |
|                                                                          |
| Processor *-------------- Change Partition Definition ---------------*   |
| Configura |                                                           |   |
| Channel S |                                                           |   |
|           | Specify or revise the following values.                   |   |
| / Partiti |                                                           |   |
| _ *       | Partition name . . . MENSA3E                              |   |
| _ *       | Partition number . . E      (same as MIF image ID)        |   |
| c *       | Partition usage  . . CF     +                             |   |
| _ *       | UID uniqueness . . . N        (Y/N)                       |   |
| ********* |                                                           | * |
|           | Description  . . . . Redbooks CF76                        |   |
|           |                                                           |   |
|           |                                                           |   |
|           |                                                           |   |
*-----------*-----------------------------------------------------------*---*
```

Figure 14-1 Processors: Change Partition Definition - Coupling Facility

14.2.2 Defining CS5 CHPIDs

When defining a CS5 CHPID to create a CF link between a CF LPAR and a z/OS LPAR, first determine which z/OS LPARs require access to which CF LPARs, how many CF links are required, and to how many different physical processors.

CS5 CF CHPIDs are defined by using Feature Code #0216(Integrated Coupling Adapter Short Reach 2.0(ICA-SR 2.0), which are two port (link) cards that are installed on the central processor complex (CPC) drawer instead of in the PCIe+ I/O drawer.

The ICA SR card has two ports (Port 1 and Port 2) that provide two physical connections to another ICA SR card on the same or different processor.

Each of the ports can have up to four CHPIDs defined to these ports.

Here are considerations for a new CS5 CHPID:

- CHPID.
- Channel path type.
- Operational mode.
- Description.
- Adapter ID (AID).

- Adapter port.
- Partition access list.
- For performance and redundancy, determine how many AID cards are installed in the processor and to what Peripheral Component Interconnect Express (PCIe) slot on what CPC drawer (for a list of installed hardware, see the physical channel ID (PCHID) / channel ID (CHID) report).

To define a CS5 CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. Press PF11, or in the CLI, enter `add` (see Figure 14-2) to add a CHPID.
5. Make the following updates and press Enter:
   – Update `Channel path ID` to `84`.
   – Update `Channel path type` to `CS5`.
   – Update `Operational mode` to `SHR`.
   – Update `Description` to the description that you want.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . 84    +            Channel ID ___  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . CS5   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . FC#0216 ICA SR 2.0 Links________                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 14-2 Processors: Add Channel Path - CS5

6. HCD prompts you to specify the adapter and port of the HCA attributes. Make the following updates (see Figure 14-3), and press Enter:
   - Update `Adapter of the HCA` to `04`.
   - Update `Port on the HCA` to `1`.

```
*------------------ Specify HCA Attributes -------------------*
|                                                              |
|                                                              |
| Specify or revise the values below.                          |
|                                                              |
| Adapter ID of the HCA  . . 04  +                             |
| Port on the HCA  . . . . . 1   +                             |
|                                                              |
|                                                              |
|                                                              |
|                                                              |
*--------------------------------------------------------------*
```

Figure 14-3 Processors: Specify HCA Attributes - CS5

7. HCD prompts you to select which partitions the CHPID should have access to. Type forward slash (/) next to the wanted partitions (see Figure 14-4), and press Enter.

```
*--------------------------- Define Access List ----------------------------*

|                                                              Row 1 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                          |
| Select one or more partitions for inclusion in the access list.          |
|                                                                          |

| Channel subsystem ID : 3                                                 |
| Channel path ID  . . : 84     Channel path type  . : CS5                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| _ 3      MENSA3A          A      OS    WTSCMXA                           |
| _ 3      MENSA3B          B      OS    RDBKPOK                           |
| _ 3      MENSA3C          C      CF    Redbooks CF86                     |
| _ 3      MENSA3D          D      CF    Redbooks CF87                     |
| / 3      MENSA3E          E      CF    Redbooks CF76                     |
| _ 3      MENSA3F          F      CF    Redbooks CF77                     |
| _ 3      MENSA31          1      OS    RDBKCRYP                          |
| / 3      MENSA32          2      OS    WTSC75                            |
| _ 3      MENSA33          3      OS    RDBKAIDV                          |
| _ 3      MENSA34          4      OS    WTSC80                            |
| _ 3      MENSA35          5      OS    WTSC81                            |
| _ 3      MENSA36          6      OS    RDBKAIWK                          |
*--------------------------------------------------------------------------*
```

Figure 14-4 Processors: Define Access List - CS5

8. Because more partitions than the selected two are defined, the Define Candidate List panel opens. For this example, we do not add any partitions to the candidate list. Press Enter. HCD returns to the Channel Path List and shows you the CHPID that was defined (see Figure 14-5).

```
                              Channel Path List      Row 1 of 7 More:      >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   N   No  FC#0216 ICA SR 2.0 Links
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
******************************* Bottom of data ********************************
```

Figure 14-5 Processors: Channel Path List - CS5

9. Proceed to define an extra CS5 CHPID as 86 to the same LPARs as `AID=10`, `Port=1` (see Figure 14-6).

```
                              Channel Path List      Row 1 of 8 More:      >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   N   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   N   No  FC#0216 ICA SR 2.0 Links
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
******************************* Bottom of data ********************************
```

Figure 14-6 Processors: Channel Path List - CS5

14.2.3 Defining a Coupling Facility link with CS5 CHPIDs

The only way to define a CS5 CHPID to another CS5 CHPID is direct connect.

Here are considerations for creating a CF link by using CS5 CHPIDs:

- The ICA SR connection is a physical cable between two ICA SR cards on the same or different processors. ICA SR 2.0 is compatible with ICA SR and ICA SR 1.1.
- Up to four logical CHPIDs per port can be defined over that physical connection.
- Determine how many CS5 CF links are required to provide enough primary and secondary links and coupling bandwidth.
- Determine which z/OS LPARs on the same or different processors need access to the CF LPARs.
- CF links provide Server Time Protocol (STP) connectivity for a Coordinated Timing Network (CTN) between processors and a sysplex.
- In this example, we connect two CS5 CHPIDs (84 and 86) on the same processor.

To define a CF link with CS5 CHPIDs, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to create the first CF link from, and press Enter.
3. Enter `s` next to the CSS ID that has the CS5 CHPID definition that you want to create the first CF link from, and press Enter.
4. Scroll through the Channel Path List until you find the first CS5 CHPID that you want to connect from, or in the CLI, enter `L 84`. In our example, we use `84`.
5. Enter `f` next to the CHPID definition (see Figure 14-7), and press Enter.

```
                              Channel Path List       Row 1 of 8 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
f 84    04/1  CS5   SHR   __  __ __   N   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   N   No  FC#0216 ICA SR 2.0 Links
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
****************************** Bottom of data *******************************
```

Figure 14-7 Processors: CF Channel Path Connectivity List - CS5

6. On the next panel, HCD prompts you to specify the second CS5 CHPID that you want to connect to. Scroll through the Channel Path List until you find the second CS5 CHPID you want to connect to, or in the CLI, enter `L 86`. In our example, we use `86`.
7. Enter `p` next to the CHPID definition (see Figure 14-8), and press Enter.

```
                     CF Channel Path Connectivity List              Row 1 of 2
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N
p 86  10/1  Y  CS5  SHR  N
******************************* Bottom of data ********************************
```

Figure 14-8 Processors: CF Channel Path Connectivity List - CS5

8. HCD prompts you to specify which CF channel path to connect to. Make the following updates (see Figure 14-9), and press Enter:
   - Update `Destination processor ID` to `MENSA`.
   - Update `Destination channel subsystem ID` to 3.
   - Update `Destination channel path ID` to 84.

```
*---------------- Connect to CF Channel Path -----------------*
|                                                              |
|                                                              |
| Specify the following values.                                |
|                                                              |
| Source processor ID  . . . . . : MENSA                       |
| Source channel subsystem ID  . : 3                           |
| Source channel path ID . . . . : 86                          |
| Source channel path type . . . : CS5                         |
|                                                              |
| Destination processor ID . . . . . . MENSA     +             |
| Destination channel subsystem ID . . 3   +                   |
| Destination channel path ID  . . . . 84  +                   |
|                                                              |
| Timing-only link . . . . . . . . . . No                      |
|                                                              |
|                                                              |
|                                                              |
|                                                              |
*--------------------------------------------------------------*
```

Figure 14-9 Processors: Connect to CF Channel Path - CS5

9. HCD checks the available CU numbers and device addresses starting at FFFF and working backwards to provide suggestions. These suggestions can be overridden or accepted. We accept the suggestions for the second CS5 CHPID (see Figure 14-10), and press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : 86        Operation mode . . : SHR     |
| Channel path type  . . . . : CS5                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FFF8                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-10 Processors: Add CF Control Unit and Devices - CS5

10. HCD provides suggested CU numbers and device addresses for the first CS5 CHPID (see Figure 14-11). Observe that the CU number is the same and that eight devices were allocated. Press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : 84        Operation mode . . : SHR     |
| Channel path type  . . . . : CS5                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FFF0                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-11 Processors: Add CF Control Unit and Devices - CS5

HCD provides a summary of the Source and Destination CF Connectivity List (see Figure 14-12).

```
                     CF Channel Path Connectivity List              Row 1 of 2
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N     MENSA.3    86  10/1  Y  CS5  SHR    CFP   8
_ 86  10/1  Y  CS5  SHR  N     MENSA.3    84  04/1  Y  CS5  SHR    CFP   8
****************************** Bottom of data ********************************
```

Figure 14-12 Processors: Add CF Channel Path Connectivity List - CS5

14.2.4 Defining CL5 and CL6 CHPIDs

When defining a CL5 or CL6 CHPID to create a CF link between a CF LPAR and a z/OS LPAR, determine which z/OS LPARs require access to which CF LPARs, how many CF links are required, and to how many different physical processors.

CL5 CF CHPIDs are defined by using Feature Code #0498 (Coupling Express3 LR 10G (CE3 LR 10G)) cards and CL6 CF CHPIDs are defined by using Feature Code #0499 (Coupling Express3 LR 25G (CE3 LR 25G)) cards, which are installed in the PCIe+ I/O drawer instead of on the CPC drawer.

The CE3 LR card has two ports (Port 1 and Port 2), which provide two physical connections between another CE LR card on the same or different processor. CHPID type CL5 can also connect to a CE LR (Feature Code #0433) on IBM z15, or a CE2 LR (Feature Code #0434) on IBM z16. However, CHPID type CL6 can only connect to a CE3 LR 25G (Feature Code #0499) on IBM z17.Each of the ports can have up to four CHPIDs defined.

Here are considerations for a new CL5 and CL6 CHPID:

- CHPID.
- CHID.
- Channel path type.
- Operational mode.
- Description.
- Partition access list.
- For performance and redundancy, determine how many CL5 and CL6 cards are installed in the processor and to what PCIe slot on what CPC drawer does the I/O cards connect to (for a list of installed hardware, see the PCHID / CHID report).

To define a CL5 CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.

4. In the CLI, enter `add` (see Figure 14-13) to add a CHPID.
5. Make the following updates and press Enter:
   - Update `Channel path ID` to `88`.
   - Update `Channel ID` to `17C`.
   - Update `Channel path type` to `CL5`.
   - Update `Operational mode` to `SHR`.
   - Update `Description` to the description that you want.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . : MENSA        Mensa                                   |
| Configuration mode . : LPAR                                                 |
| Channel Subsystem ID : 3                                                    |
|                                                                             |
| Channel path ID  . . . . 88    +            Channel ID  17C  +              |
| Number of CHPIDs . . . . 1                                                  |
| Channel path type  . . . CL5   +                                            |
| Operation mode . . . . . SHR   +                                            |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +        |
| Description  . . . . . . FC#0498 Coupling Express3 LR 10G                   |
|                                                                             |
| Specify the following values only if connected to a switch:                 |
| Dynamic entry switch ID  __  + (00 - FF)                                    |
| Entry switch ID  . . . . __  +                                              |
| Entry port . . . . . . . __  +                                              |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 14-13 Processors: Add Channel Path - CL5

6. HCD prompts you to specify the coupling PCHID/Port attributes. Update `Coupling Port` to `1` (see Figure 14-14), and press Enter.

```
*----------- Specify Coupling Pchid/Port Attribute -----------*
|                                                              |
|                                                              |
| Specify or revise the values below.                          |
|                                                              |
| Physical channel ID  . . . 17C                               |
| Coupling port  . . . . . . 1   +                             |
|                                                              |
|                                                              |
|                                                              |
|                                                              |
*--------------------------------------------------------------*
```

Figure 14-14 Processors: Specify Coupling PCHIDs/Port Attributes - CL5

7. HCD prompts you to select which partitions the CHPID should have access to. Enter forward slash (/) next to the partitions that you want (see Figure 14-15), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                             Row 1 of 15 |
| Command ===> ________________________________________ Scroll ===> CSR    |
|                                                                          |
| Select one or more partitions for inclusion in the access list.         |
|                                                                          |
| Channel subsystem ID : 3                                                 |
| Channel path ID  . . : 88     Channel path type  . : CL5                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| _ 3      MENSA3A          A      OS    WTSCMXA                           |
| _ 3      MENSA3B          B      OS    RDBKPOK                           |
| _ 3      MENSA3C          C      CF    Redbooks CF86                     |
| _ 3      MENSA3D          D      CF    Redbooks CF87                     |
| / 3      MENSA3E          E      CF    Redbooks CF76                     |
| _ 3      MENSA3F          F      CF    Redbooks CF77                     |
| _ 3      MENSA31          1      OS    RDBKCRYP                          |
| / 3      MENSA32          2      OS    WTSC75                            |
| _ 3      MENSA33          3      OS    RDBKAIDV                          |
| _ 3      MENSA34          4      OS    WTSC80                            |
| _ 3      MENSA35          5      OS    WTSC81                            |
| _ 3      MENSA36          6      OS    RDBKAIWK                          |
*--------------------------------------------------------------------------*
```

Figure 14-15 Processors: Define Access List - CL5

8. Because more partitions than the selected two are defined, the Define Candidate List panel opens. For this example, we do not add any partitions in the access list. Press Enter. HCD returns to the Channel Path List and shows you the CHPID that was defined (see Figure 14-16).

```
                              Channel Path List       Row 1 of 9 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 88    17C/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
******************************* Bottom of data ********************************
```

Figure 14-16 Processors: Channel Path List - CL5

9. Define an extra CL5 CHPID as `8A` to the same LPARs as `CHID=198, Port=1` (see Figure 14-17).

```
                              Channel Path List      Row 1 of 10 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 88    17C/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8A    198/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
```

Figure 14-17 Processors: Channel Path List - CL5

CL6 CHPIDs can be defined using similar steps. To define a CL6 CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. In the CLI, enter `add` (see Figure 14-18) to add a CHPID.
5. Make the following updates and press Enter:
   - Update `Channel path ID` to `8C`.
   - Update `Channel ID` to `1F8`.
   - Update `Channel path type` to `CL6`.
   - Update `Operational mode` to `SHR`.
   - Update `Description` to the description that you want.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . 8C    +           Channel ID  1F8  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . CL6   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0499 Coupling Express3 LR 25G                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 14-18 Processors: Add Channel Path - CL6

6. HCD prompts you to specify the coupling PCHID/Port attributes. Update `Coupling Port` to `1` (see Figure 14-19), and press Enter.

```
*----------- Specify Coupling Pchid/Port Attribute -----------*
|                                                            |
|                                                            |
| Specify or revise the values below.                        |
|                                                            |
| Physical channel ID  . . . 1F8                             |
| Coupling port  . . . . . . 1   +                           |
|                                                            |
|                                                            |
|                                                            |
|                                                            |
*------------------------------------------------------------*
```

Figure 14-19 Processors: Specify Coupling PCHIDs/Port Attributes - CL6

7. HCD prompts you to select which partitions the CHPID should have access to. Enter forward slash (/) next to the partitions that you want (see Figure 14-20), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                               Row 1 of 15 |
| Command ===> ________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.          |
|                                                                           |
| Channel subsystem ID : 3                                                  |
| Channel path ID  . . : 8C     Channel path type  . : CL6                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| _ 3      MENSA3A          A      OS    WTSCMXA                            |
| _ 3      MENSA3B          B      OS    RDBKPOK                            |
| _ 3      MENSA3C          C      CF    Redbooks CF86                      |
| / 3      MENSA3D          D      CF    Redbooks CF87                      |
| _ 3      MENSA3E          E      CF    Redbooks CF76                      |
| _ 3      MENSA3F          F      CF    Redbooks CF77                      |
| _ 3      MENSA31          1      OS    RDBKCRYP                           |
| _ 3      MENSA32          2      OS    WTSC75                             |
| _ 3      MENSA33          3      OS    RDBKAIDV                           |
| / 3      MENSA34          4      OS    WTSC80                             |
| _ 3      MENSA35          5      OS    WTSC81                             |
| _ 3      MENSA36          6      OS    RDBKAIWK                           |
*---------------------------------------------------------------------------*
```

Figure 14-20 Processors: Define Access List - CL6

8. Because more partitions than the selected two are defined, the Define Candidate List panel opens. For this example, we do not add any partitions in the access list. Press Enter. HCD returns to the Channel Path List and shows you the CHPID that was defined (see Figure 14-21).

```
                              Channel Path List       Row 1 of 11 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 88    17C/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8A    198/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __   N   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
```

Figure 14-21 Processors: Channel Path List - CL6

9. Define an extra CL6 CHPID as `8E` to the same LPARs as `CHID=21C, Port=1` (see Figure 14-22).

```
                              Channel Path List      Row 1 of 12 More:       >
Command ===> ______________________________________________ Scroll ===> CSR


Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 88    17C/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8A    198/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __   N   No  FC#0499 Coupling Express3 LR 25G
_ 8E    21C/1 CL6   SHR   __  __ __   N   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
```

Figure 14-22 Processors: Channel Path List - CL6

14.2.5 Defining a Coupling Facility link with CL5 or CL6 CHPIDs

The only way to define a CL5 CHPID to another CL5 CHPID or a CL6 CHPID to another CL6 CHPID is direct connect.

Here are considerations for creating a CF link by using CL5 or CL6 CHPIDs:

- The CE LR connection is a physical cable between two CE LR cards. A CL5 CHPID can connect to another CL5 CHPID on IBM z15, IBM z16, or IBM z17, but a CL6 CHPID can only connect to another CL6 CHPID on IBM z17 machines.
- Up to four logical CHPIDs per port can be defined over that physical connection.
- Determine how many CL5 and CL6 CF links are required to provide enough primary and secondary links and coupling bandwidth.
- Determine which z/OS LPARs on the same or different processors need access to the CF LPARs.
- CF links also provide STP connectivity for a CTN between processors and a sysplex.
- In this example, we connect two CL5 CHPIDs (88 and 8A) on the same processor.

To define a CF link with CL5 CHPIDs, complete the following steps. You can define a CF link with CL6 CHPIDs with the same steps.

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to create the first CF link from, and press Enter.

3. Enter `s` next to the CSS ID that has the CL5 CHPID definition that you want to create the first CF link from, and press Enter.
4. Scroll through the Channel Path List until you find the first CL5 CHPID that you want to connect from, or in the CLI, enter `L 88`. In our example, we use `88`.
5. Enter `f` next to the CHPID definition (see Figure 14-23), and press Enter.

```
                              Channel Path List      Row 1 of 12 More:       >
Command ===> _______________________________________________ Scroll ===> CSR


Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 54    1A0   FC    SPAN  01  __ __       No  FCTC 32G LX
_ 55    26C   FC    SPAN  02  __ __       No  FCTC 32G LX
_ 84    04/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
_ 86    10/1  CS5   SHR   __  __ __   Y   No  FC#0216 ICA SR 2.0 Links
f 88    17C/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8A    198/1 CL5   SHR   __  __ __   N   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __   N   No  FC#0499 Coupling Express3 LR 25G
_ 8E    21C/1 CL6   SHR   __  __ __   N   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
```

Figure 14-23 Processors: CF Channel Path Connectivity List - CL5

6. HCD prompts you to specify the second CL5 CHPID that you want to connect to. Scroll through the Channel Path List until you find the second CL5 CHPID that you want to connect to, or in the CLI, enter `L 8A`. In our example, we use `8A`.
7. Type `p` next to the CHPID definition (see Figure 14-24), and press Enter.

```
---------------------------------------------------------------------------
                     CF Channel Path Connectivity List            Row 1 of 6
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N     MENSA.3    86  10/1  Y  CS5  SHR    CFP   8
_ 86  10/1  Y  CS5  SHR  N     MENSA.3    84  04/1  Y  CS5  SHR    CFP   8
_ 88  17C/1 Y  CL5  SHR  N
p 8A  198/1 Y  CL5  SHR  N
_ 8C  1F8/1 Y  CL6  SHR  N
_ 8E  21C/1 Y  CL6  SHR  N
******************************* Bottom of data ********************************
```

Figure 14-24 Processors: CF Channel Path Connectivity List - CL5

8. HCD prompts you to specify which CF channel path to connect to. Make the following updates (see Figure 14-25), and press Enter:
   - Update `Destination processor ID` to `MENSA`.
   - Update `Destination channel subsystem ID` to 3.
   - Update `Destination channel path ID` to 88.

```
*---------------- Connect to CF Channel Path -----------------*
|                                                              |
|                                                              |
| Specify the following values.                                |
|                                                              |
| Source processor ID  . . . . . : MENSA                       |
| Source channel subsystem ID  . : 3                           |
| Source channel path ID . . . . : 8A                          |
| Source channel path type . . . : CL5                         |
|                                                              |
| Destination processor ID . . . . . . MENSA     +             |
| Destination channel subsystem ID . . 3   +                   |
| Destination channel path ID  . . . . 88  +                   |
|                                                              |
| Timing-only link . . . . . . . . . . No                      |
|                                                              |
|                                                              |
|                                                              |
|                                                              |
*--------------------------------------------------------------*
```

Figure 14-25 Processors: Connect to CF Channel Path - CL5

9. HCD checks the available CU numbers and device addresses starting at FFFF and works backwards to provide suggestions. These suggestions can be overridden or accepted. We accept the suggestions for the second CL5 CHPID (see Figure 14-26), and press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : 8A        Operation mode . . : SHR     |
| Channel path type  . . . . : CL5                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FF88                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-26 Processors: Add CF Control Unit and Devices - CL5

10. HCD provides suggested CU numbers and device addresses for the first CL5 CHPID (see Figure 14-27). Observe that the CU number is the same and that eight devices were allocated. Press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : 88        Operation mode . . : SHR     |
| Channel path type  . . . . : CL5                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FF80                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-27 Processors: Add CF Control Unit and Devices - CL5

HCD provides a summary of the Source and Destination CF Connectivity List (see Figure 14-28).

```
                     CF Channel Path Connectivity List            Row 1 of 6
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N     MENSA.3    86  10/1  Y  CS5  SHR    CFP   8
_ 86  10/1  Y  CS5  SHR  N     MENSA.3    84  04/1  Y  CS5  SHR    CFP   8
_ 88  17C/1 Y  CL5  SHR  N     MENSA.3    8A  198/1 Y  CL5  SHR    CFP   8
_ 8A  198/1 Y  CL5  SHR  N     MENSA.3    88  17C/1 Y  CL5  SHR    CFP   8
_ 8C  1F8/1 Y  CL6  SHR  N
_ 8E  21C/1 Y  CL6  SHR  N
****************************** Bottom of data *******************************
```

Figure 14-28 Processors: Add CF Channel Path Connectivity List - CL5

A CF link with CL6 CHPIDs can be connected with the same steps. Unlike CL5 CHPIDs, CL6 CHPIDs default to a definition of 32 subchannels/devices per CHPID, with an option for 8. After Connecting the CL6 CHPIDs, you can confirm the Source and Destination information in the HCD CF Connectivity List (see Figure 14-29).

```
                     CF Channel Path Connectivity List            Row 1 of 6
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N     MENSA.3    86  10/1  Y  CS5  SHR    CFP   8
_ 86  10/1  Y  CS5  SHR  N     MENSA.3    84  04/1  Y  CS5  SHR    CFP   8
_ 88  17C/1 Y  CL5  SHR  N     MENSA.3    8A  198/1 Y  CL5  SHR    CFP   8
_ 8A  198/1 Y  CL5  SHR  N     MENSA.3    88  17C/1 Y  CL5  SHR    CFP   8
_ 8C  1F8/1 Y  CL6  SHR  N     MENSA.3    8E  21C/1 Y  CL6  SHR    CFP   32
_ 8E  21C/1 Y  CL6  SHR  N     MENSA.3    8C  1F8/1 Y  CL6  SHR    CFP   32
****************************** Bottom of data *******************************
```

Figure 14-29 Processors: Add CF Channel Path Connectivity List - CL5 and CL6

14.2.6 Defining ICP CHPIDs

When defining an ICP CHPID to create a CF link between a CF LPAR and a z/OS LPAR, determine which z/OS LPARs require access to which CF LPARs, and how many CF links are required within the same physical processor.

ICP CF CHPIDs are defined logically and internally to the processor and require no installed hardware.

> **Note:** The maximum number of ICP CHPIDs for an IBM z17 ME1 is 64.

Each of the logical ICP links can support only one CHPID at each end of the link. However, the CHPIDS can be spanned across multiple CSSs.

Here are considerations for a new ICP CHPID:

- CHPID
- Channel path type
- Operational mode
- Description
- Partition access list

To define an ICP CHPID and provide access to a partition, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to add a CHPID to, and press Enter.
3. Enter `s` next to the CSS ID that you want to add a CHPID to, and press Enter.
4. In the CLI, enter `add` (see Figure 14-30) to add a CHPID.
5. Make the following updates, and press Enter:
   - Update `Channel path ID` to `FC`.
   - Update `Channel path type` to `ICP`.
   - Update `Operational mode` to `SHR`.
   - Update `Description` to the description that you want.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MENSA        Mensa                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 3                                                  |
|                                                                           |
| Channel path ID  . . . . FC    +            Channel ID ___  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . ICP   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . ICP Internal Coupling Peer______                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 14-30 Processors: Add Channel Path - ICP

6. HCD prompts you to select which partitions the CHPID should have access to. Enter a forward slash (/) next to the partitions that you want (see Figure 14-31), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                              Row 1 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                          |
| Select one or more partitions for inclusion in the access list.          |
|                                                                          |
| Channel Subsystem ID : 3                                                 |
| Channel path ID  . . : FC     Channel path type  . : ICP                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |

| _ 3      MENSA3A          A      OS    WTSCMXA                           |
| _ 3      MENSA3B          B      OS    RDBKPOK                           |
| _ 3      MENSA3C          C      CF    Redbooks CF86                     |
| _ 3      MENSA3D          D      CF    Redbooks CF87                     |
| / 3      MENSA3E          E      CF    Redbooks CF76                     |
| _ 3      MENSA3F          F      CF    Redbooks CF77                     |
| _ 3      MENSA31          1      OS    RDBKCRYP                          |
| / 3      MENSA32          2      OS    WTSC75                            |
| _ 3      MENSA33          3      OS    RDBKAIDV                          |
| _ 3      MENSA34          4      OS    WTSC80                            |
| _ 3      MENSA35          5      OS    WTSC81                            |
| _ 3      MENSA36          6      OS    RDBKAIWK                          |
*--------------------------------------------------------------------------*
```

Figure 14-31 Processors: Define Access List - ICP

Because more partitions than the selected two are defined, the Define Candidate List panel opens. For this example, we do not add any partitions in the access list. Press Enter.

HCD returns to the Channel Path List and shows you the CHPID that was defined (see Figure 14-32).

```
                              Channel Path List      Row 5 of 13 More:       >
Command ===> ______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 88    17C/1 CL5   SHR   __  __ __    Y   No  FC#0498 Coupling Express3 LR 10G
_ 8A    198/1 CL5   SHR   __  __ __    Y   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __    Y   No  FC#0499 Coupling Express3 LR 25G
_ 8E    21C/1 CL6   SHR   __  __ __    Y   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __        No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __        No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __        No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __        No  NetExpSR 10G IBM Redbooks 129.
_ FC    _____ ICP   SHR   __  __ __    N   No  ICP Internal Coupling Peer
****************************** Bottom of data *******************************
```

Figure 14-32 Processors: Channel Path List - ICP

7. Define an extra ICP CHPID as FD to the same LPARs (see Figure 14-33).

```
                              Channel Path List      Row 6 of 14 More:       >
Command ===> ______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 8A    198/1 CL5   SHR   __  __ __    Y   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __    Y   No  FC#0499 Coupling Express3 LR 25G
_ 8E    21C/1 CL6   SHR   __  __ __    Y   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __        No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __        No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __        No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __        No  NetExpSR 10G IBM Redbooks 129.
_ FC    _____ ICP   SHR   __  __ __    N   No  ICP Internal Coupling Peer
_ FD    _____ ICP   SHR   __  __ __    N   No  ICP Internal Coupling Peer
```

Figure 14-33 Processors: Channel Path List - ICP

14.2.7 Defining a Coupling Facility link with ICP CHPIDs

The only way to define an ICP CHPID to another ICP CHPID is as a logical internal connection within the processor.

Here are considerations for creating a CF link by using ICP CHPIDs:

- Each of the logical ICP links can support only one CHPID at each end of the link.
- Determine how many ICP CF links are required to provide enough coupling bandwidth.
- Determine which z/OS LPARs on the same processors need access to the CF LPARs.
- In this example, we connect two ICP CHPIDs (Fibre Channel (FC) and FD) on the same processor.

To define a CF link with ICP CHPIDs, complete the following steps:

1. From the main HCD panel, select option `1.3. Processors`.
2. Enter `s` next to the processor that you want to create the first CF link from, and press Enter.
3. Enter `s` next to the CSS ID that has the ICP CHPID definition that you want to create the first CF link from, and press Enter.
4. Scroll through the Channel Path List until you find the first ICP CHPID that you want to connect from, or in the CLI, enter `L FC`. In our example, we use `FC`.

5. Enter `f` next to the CHPID definition (see Figure 14-34), and press Enter.

```
                              Channel Path List      Row 6 of 14 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . . : MENSA        Mensa
Configuration mode . : LPAR
Channel Subsystem ID : 3

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 8A    198/1 CL5   SHR   __  __ __   Y   No  FC#0498 Coupling Express3 LR 10G
_ 8C    1F8/1 CL6   SHR   __  __ __   Y   No  FC#0499 Coupling Express3 LR 25G
_ 8E    21C/1 CL6   SHR   __  __ __   Y   No  FC#0499 Coupling Express3 LR 25G
_ 94    19C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 95    1B0   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
_ 96    25C   OSH   SPAN  __  __ __       No  NetExpSR 10G IBM Redbooks 9.
_ 97    278   OSH   SHR   __  __ __       No  NetExpSR 10G IBM Redbooks 129.
f FC    _____ ICP   SHR   __  __ __   N   No  ICP Internal Coupling Peer
_ FD    _____ ICP   SHR   __  __ __   N   No  ICP Internal Coupling Peer
```

Figure 14-34 Processors: CF Channel Path Connectivity List - ICP

6. HCD prompts you to specify the second ICP CHPID that you want to connect to. Scroll through the Channel Path List until you find the second ICP CHPID that you want to connect to, or in the CLI, enter `L FD`. In our example, we use `FD`.
7. Type `p` next to the CHPID definition (see Figure 14-35), and press Enter.

```
---------------------------------------------------------------------------
                     CF Channel Path Connectivity List              Row 1 of 8
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . . : MENSA     Mensa
Source channel subsystem ID . . : 3
Source partition name . . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 84  04/1  Y  CS5  SHR  N     MENSA.3    86  10/1  Y  CS5  SHR    CFP   8
_ 86  10/1  Y  CS5  SHR  N     MENSA.3    84  04/1  Y  CS5  SHR    CFP   8
_ 88  17C/1 Y  CL5  SHR  N     MENSA.3    8A  198/1 Y  CL5  SHR    CFP   8
_ 8A  198/1 Y  CL5  SHR  N     MENSA.3    88  17C/1 Y  CL5  SHR    CFP   8
_ 8C  1F8/1 Y  CL6  SHR  N     MENSA.3    8E  21C/1 Y  CL6  SHR    CFP   32
_ 8E  21C/1 Y  CL6  SHR  N     MENSA.3    8C  1F8/1 Y  CL6  SHR    CFP   32
_ FC        Y  ICP  SHR  N
p FD        Y  ICP  SHR  N
```

Figure 14-35 Processors: CF Channel Path Connectivity List - ICP

8. HCD prompts you to specify which CF channel path to connect to. Make the following updates (see Figure 14-36), and press Enter:
   - Update `Destination processor ID` to `MENSA`.
   - Update `Destination channel subsystem ID` to 3.
   - Update `Destination channel path ID` to `FC`.

```
*---------------- Connect to CF Channel Path -----------------*
|                                                              |
|                                                              |
| Specify the following values.                                |
|                                                              |
| Source processor ID  . . . . . : MENSA                       |
| Source channel subsystem ID  . : 3                           |
| Source channel path ID . . . . : FD                          |
| Source channel path type . . . : ICP                         |
|                                                              |
| Destination processor ID . . . . . . MENSA     +             |
| Destination channel subsystem ID . . 3   +                   |
| Destination channel path ID  . . . . FC  +                   |
|                                                              |
| Timing-only link . . . . . . . . . . No                      |
|                                                              |
|                                                              |
|                                                              |
|                                                              |
*--------------------------------------------------------------*
```

Figure 14-36 Processors: Connect to CF Channel Path - ICP

9. HCD checks the available CU numbers and device addresses starting at FFFF and works backwards to provide suggestions. These suggestions can be overridden or accepted. On IBM z17, ICP CHPIDs can be defined with 8 subchannels/devices per CHPID, although configuration with 7 subchannels/devices is allowed and supported by HCD. We accept the suggestions for the second ICP CHPID (see Figure 14-37), and press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : FD         Operation mode . . : SHR    |
| Channel path type  . . . . : ICP                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FF78                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-37 Processors: Add CF Control Unit and Devices - ICP

10. HCD provides suggested CU numbers and device addresses for the first ICP CHPID (see Figure 14-38). Observe that the CU number is the same and that seven devices were allocated. Press Enter.

```
*----------------- Add CF Control Unit and Devices -------------------*
|                                                                     |
|                                                                     |
| Confirm or revise the CF control unit number and device numbers     |
| for the CF control unit and devices to be defined.                  |
|                                                                     |
| Processor ID . . . . . . . : MENSA                                  |
| Channel subsystem ID . . . : 3                                      |
| Channel path ID  . . . . . : FC        Operation mode . . : SHR     |
| Channel path type  . . . . : ICP                                    |
|                                                                     |
| Control unit number  . . . . FFFE  +                                |
|                                                                     |
| Device number  . . . . . . . FF70                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-38 Processors: Add CF Control Unit and Devices - ICP

HCD provides a summary of the Source and Destination CF Connectivity List (see Figure 14-39).

```
                     CF Channel Path Connectivity List              Row 7 of 8
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MENSA     Mensa
Source channel subsystem ID . : 3
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ FC        Y  ICP  SHR  N     MENSA.3    FD        Y  ICP  SHR    CFP   8
_ FD        Y  ICP  SHR  N     MENSA.3    FC        Y  ICP  SHR    CFP   8
******************************* Bottom of data ********************************
```

Figure 14-39 Processors: Add CF Channel Path Connectivity List - ICP

15

Adding Peripheral Component Interconnect Express devices

This chapter describes the steps to define Peripheral Component Interconnect Express (PCIe) functions, features, and devices for Internal Shared Memory (ISM), Network Express (NETH), and zHyperLink.

It includes a list of these potential configuration items and a short description about how to do each of them by using Hardware Configuration Definition (HCD). It also includes an overview of PCIe functions.

> **Note:** The examples that are shown in this chapter are based on the IBM z17 ME1 (9175). Throughout this chapter, we might refer to this machine as IBM z17.
>
> **Note:** Not all the following configuration items are necessarily required for your installation. In addition, the examples that are presented are not exhaustive.
>
> **Note:** All definitions and commands in the examples of this book were done in z/OS 3.1. All definitions and values are examples only.

This chapter includes the following topics:

- Defining PCIe functions by using HCD
- PCIe feature definitions

15.1 Defining PCIe functions by using HCD

When defining new I/O components in an input/output definition file (IODF), certain definitions like operating system configurations (OSCONFIGs), partitions, Fibre Connection (FICON) switches, control units (CUs), and devices must be done first. After these items are defined, then the following definitions can be made:

- Defining an ISM PCIe function that is associated with Hypersocket (IQD), OSA Express (OSD) or Network Express (OSH) for SMC-D virtual network connection (internal).
- Defining a NETH PCIe function that is associated with a Network Express adapter for SMC-R communication.
- Defining a HYL (zHyperLink) PCIe function.

> **Note:** Starting with the IBM z15, the IBM zEnterprise Data Compression (zEDC) PCIe feature is no longer available as a separate hardware feature. Compression functions are implemented directly on the processor chip. Therefore, an HCD configuration is no longer required.
>
> **Note:** IBM z17 no longer supports RoCE Express adapters and the PCIe ROC and ROC2 functions. SMC-R communication is enabled by associating a Network Express adapter port with a NETH PCIe function.

15.2 PCIe feature definitions

This section provides a brief overview of the `FUNCTION` statement and covers the HCD definitions of ISM, Network Express (NETH), and zHyperLink (HYL) PCIe features.

15.2.1 Overview

PCIe adapters that are attached to a system can provide the operating system (OS) with various “PCIe functions” that are used by entitled logical partitions (LPARs).

Currently, HCD supports the following features for an IBM z17:

- **ISM** PCIe Adapter: A virtual PCIe (vPCIe) adapter for SMC-D communications, for which a virtual channel ID (VCHID) must be defined.
- **Network Express:** PCIe functions of type NETH for SMC-R may be assigned to external physical networks by specifying corresponding PNETIDs.
- **zHyperLink**: Requires a PCIe function of type HYL with an attribute for identifying a port on the adapter to which the function is related.

HCD provides dialog boxes to define, change, delete, and view PCIe functions, and to control which LPARs access which PCIe functions.

In addition, HCD provides the following reports:

- The PCIe Function Summary Report displays the partitions in the access and candidate lists, which are entitled to access the available PCIe functions.
- The PCIe Function Compare Report shows the changes of PCIe functions between processors of two IODFs.

IOCP supports the I/O configuration statement `FUNCTION` for defining and configuring PCIe functions. To define PCIe functions for the I/O subsystem, you must specify at least the following parameters:

- **FID**: The function number that identifies the function. This number is a 4 digit hexadecimal arbitrary number and is unique for the entire CPC. On IBM z17 the maximum FID is 4FFF. An adapter can have multiple FIDs associated with it.
- **CHID**: Specifies either a physical channel (PCHID) or a virtual channel (VCHID) associated with the function.
- **PART**: Specifies the availability of FIDs to logical partitions. There is at most one partition designated as the *access* LPAR, and a group of LPARs specified as the *candidate list*. The access LPAR is the one that has the channel path online after a POR. Partitions in the candidate list can access the function through reconfiguration commands. Functions can not be shared by multiple LPARs at the same time.
- **VF**: Specifies a virtual function so that multiple logical partitions can use the same CHIDs (PCHID/port or VCHID) associated with de function.
- **TYPE**: The type of the virtual function. In this book we cover the **ISM** (SMC-D), **NETH** (SMC-R) and the **HYL** (zHyperLink) function types.
- **PNETID**: Up to four physical network identifiers that are used for the ISM and NETH functions.

> **Note:** The support of virtual functions (VFs), the allowed range of virtual function IDs (VFIDs), and the support of PNETIDs depends on the processor type and support level. For more information, see *Input/Output Configuration Program User's Guide,* SB10-7183. HCD offers prompts for VFIDs and ensures that the validation rules are fulfilled.

In the following configurations we are going to see how these definitions are made in HCD panels for the functions that are covered in this book. Example 15-1 shows definitions for a zHyperLink card for Port 1 and multiple VFIDs, each assigned to a specific LPAR.

Example 15-1 zHyperLink definitions

```
RESOURCE PARTITION=((CSS(0),(MENSA0A,A),(MENSA0B,B),(MENSA0C,C*
      ),(MENSA0D,D),(MENSA01,1),(MENSA02,2),(MENSA03,3),(MENSA*
      04,4),(MENSA05,5),(MENSA06,6),(MENSA07,7),(MENSA08,8),(M*
      ENSA09,9),(*,E),(*,F)),(CSS(1),(MENSA11,1),(MENSA12,2),(*
      MENSA13,3),(MENSA14,4),(MENSA15,5),(MENSA16,6),(MENSA17,*
      7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS*
      (2),(MENSA21,1),(MENSA22,2),(MENSA23,3),(MENSA24,4),(MEN*
      SA25,5),(MENSA26,6),(MENSA27,7),(MENSA28,8),(MENSA29,9),*
      (*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(3),(MENSA3A,A)*
      ,(MENSA3B,B),(MENSA3C,C),(MENSA3D,D),(MENSA3E,E),(MENSA3*
      F,F),(MENSA31,1),(MENSA32,2),(MENSA33,3),(MENSA34,4),(ME*
      NSA35,5),(MENSA36,6),(MENSA37,7),(MENSA38,8),(MENSA39,9)*
      ),(CSS(4),(MENSA4A,A),(MENSA4B,B),(MENSA41,1),(MENSA42,2*
      ),(MENSA43,3),(MENSA44,4),(MENSA45,5),(MENSA46,6),(MENSA*
      47,7),(MENSA48,8),(MENSA49,9),(*,C),(*,D),(*,E),(*,F)),(*
      CSS(5),(MENSA51,1),(MENSA52,2),(*,3),(*,4),(*,5),(*,6),(*
      *,7),(*,8),(*,9),(*,A)))
*********************************************************************
```

```
* FOR zHyperLink
**********************************************************************
FUNCTION FID=2102,VF=1,PCHID=1FC,PART=((MENSA02),(=)),        *
      TYPE=HYL,PORT=1
FUNCTION FID=2108,VF=2,PCHID=1FC,PART=((MENSA08),(=)),        *
      TYPE=HYL,PORT=1
FUNCTION FID=2132,VF=3,PCHID=1FC,PART=((MENSA32),(=)),        *
      TYPE=HYL,PORT=1
**
```

Table 15-1 shows the list of applicable functions to the various function types.

Table 15-1 Keyword applicability for functions

Function	**Function type**		
	HYL (zHyperLink)	**ISM (SMC-D)**	**NETH (SMC-R)**
Physical channel ID (PCHID)	Yes	No	Yes
Virtual channel ID (VCHID)	No	Yes	No
PNETID	No	Yes	Yes
PORT	Yes	No	No
VF	Yes	Yes	Yes

For more information about the maximum values for each machine type, see *IBM Z Input/Output Configuration Program User's Guide*, SB10-7183.

15.2.2 Defining an ISM PCIe function

The IBM z17 system supports ISM virtual PCIe devices to enable optimized cross-LPAR TCP communications by using socket-based Direct Memory Access (DMA), that is, Shared Memory Communication - Direct (SMC-D) over Internal Shared Memory (ISM).

ISM Overview

ISM is defined as PCIe devices. VCHIDs and FIDs/VFs must be defined in HCD. A VCHID represents a virtual PCIe adapter, which also represents a unique isolated ISM network with a unique PNETID.

IBM z17 supports up to 32 ISM VCHIDs per CPC and up to 255 VFs per VCHID. Each VF represents a TCP/IP stack that can communicate over the same ISM VCHID with other LPARs in the same CPC. On IBM z17 the VCHID is a 3 digit hexadecimal number in the range 7C0 to 7FF.

For more information about the management of SMC-D, see *z/OS Communications Server: IP Configuration Guide*.

ISM Configuration

In the following example we do the PCIe definition for an ISM network in VCHID 7F1 and PNETID PERFNET. The function definitions are:

- `CHID=7F1` to FID 0040 (`VF=1`) for LPAR MENSA32 on `CPC=MENSA`
- `CHID=7F1` to FID 0041 (`VF=2`) for LPAR MENSA33 on `CPC=MENSA`

Figure 15-1 shows a diagram for our configuration.

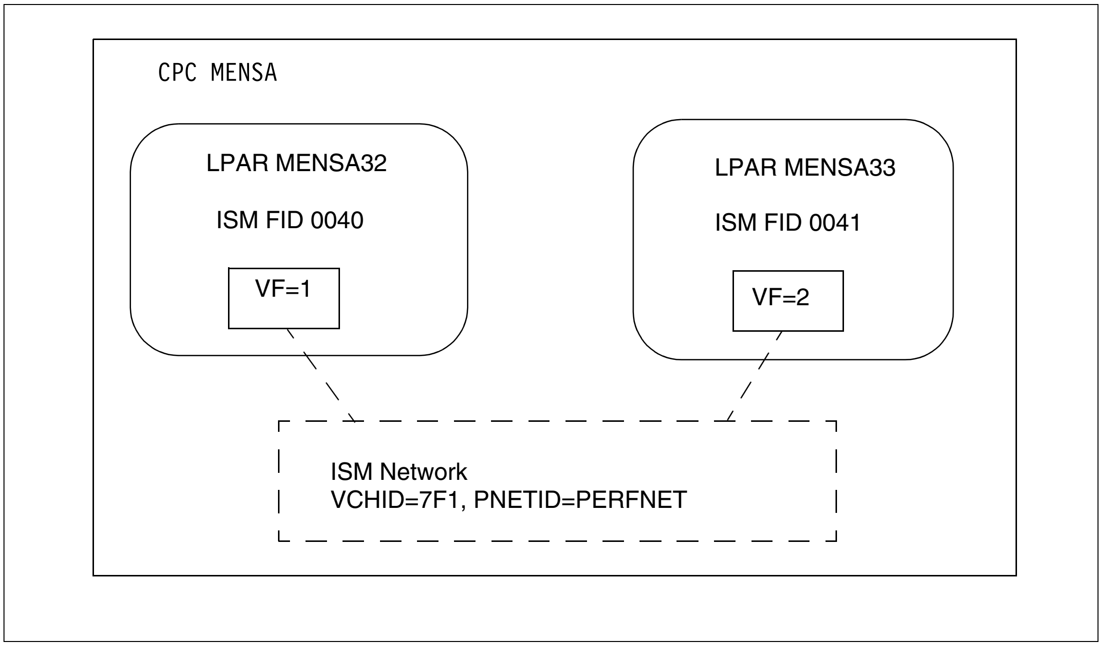

Figure 15-1 ISM Network for Sample Definitions

Complete the following steps:

1. From the main HCD panel, select option `1.3. Processor List`.
2. Enter `f` (work with PCIe functions) next to the processor (`MENSA`) to which you want to define the ISM functions, as shown in Figure 15-2, and press Enter.

```
------------------------------------------------------------------------------
                                Processor List       Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
f MENSA    9175     ME1      LPAR  0B9FB89175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
******************************* Bottom of data ********************************
```

Figure 15-2 Processor List: Work with PCIe functions - ISM

3. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (Figure 15-3).

```
--------------------------------------------------------------------------------
                              PCIe Function List     Row 1 of 22 More:       >
Command ===> add____________________________________________ Scroll ===> PAGE

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
```

Figure 15-3 PCIe Function List: Add line command - ISM

4. Make the following updates (Figure 15-4), and press Enter:
   - Update `Function ID` to `0040`.
   - Update `Type` to `ISM`.
   - Update `Channel ID` to `7F1`.
   - Update `Virtual Function ID` to `1`.
   - Update `Description` to the description that you want.

```
............................. Add PCIe Function .............................
.                                                                             .
.                                                                             .
. Specify or revise the following values.                                     .
.                                                                             .
. Processor ID . . . . . . . : MENSA        Mensa                             .
.                                                                             .
. Function ID  . . . . . . . . 0040                                           .
. Type . . . . . . . . . . . . ISM_________  +                                .
.                                                                             .
. Channel ID . . . . . . . . . 7F1  +                                         .
. Port . . . . . . . . . . . . _    +                                         .
. Virtual Function ID  . . . . 1__  +     Physical Function    _  (Y/N)       .
. Number of virtual functions  1                                              .
. UID  . . . . . . . . . . . . ____                                           .
.                                                                             .
. Description  . . . . . . . . ISM FID=0040 VFID=1 CHID=7F1____               .
.                                                                             .
.                                                                             .
.                                                                             .
...............................................................................
```

Figure 15-4 PCIe Function List: Add PCIe function - ISM

5. Update `Physical network ID` 1 to `PERFNET` in the Add/Modify Physical Network IDs panel (Figure 15-5), and press Enter.

```
····················· Add/Modify Physical Network IDs ······················
·                                                                         ·
·                                                                         ·
· If the Channel ID (CHID) is associated to one or more physical networks, ·
· specify each physical network ID corresponding to each applicable        ·
· physical port.                                                          ·
·                                                                         ·
· Physical network ID 1  . . PERFNET_________                             ·
· Physical network ID 2  . . ________________                             ·
· Physical network ID 3  . . ________________                             ·
· Physical network ID 4  . . ________________                             ·
·                                                                         ·
·                                                                         ·
·                                                                         ·
·  F1=Help     F2=Split    F3=Exit     F5=Reset    F9=Swap    F12=Cancel  ·
···························································································
```

Figure 15-5 Add PCIe Function: Add/Modify Physical Network IDs - ISM

6. Select the access LPAR for the function. In our example, we use LPAR `MENSA32(OS)`. Press Enter (Figure 15-6).

```
··························· Define Access List ····························
·                                                          Row 29 of 57 ·
· Command ===> _________________________________________ Scroll ===> CSR  ·
·                                                                         ·
· Select one partition for the access list.                               ·
·                                                                         ·
· Function ID  . . . . : 0040                                             ·
·                                                                         ·
· / CSS ID Partition Name   Number Usage Description                      ·
· _ 2      MENSA29          9      OS                                     ·
· _ 3      MENSA3A          A      OS    WTSCMXA                          ·
· _ 3      MENSA3B          B      OS    RDBKPOK                          ·
· _ 3      MENSA3C          C      CF    Redbooks CF86                    ·
· _ 3      MENSA3D          D      CF    Redbooks CF87                    ·
· _ 3      MENSA3E          E      CF    Redbooks CF76                    ·
· _ 3      MENSA3F          F      CF    Redbooks CF77                    ·
· _ 3      MENSA31          1      OS    RDBKCRYP                         ·
· / 3      MENSA32          2      OS    WTSC75                           ·
· _ 3      MENSA33          3      OS    RDBKAIDV                         ·
· _ 3      MENSA34          4      OS    WTSC80                           ·
· _ 3      MENSA35          5      OS    WTSC81                           ·
· _ 3      MENSA36          6      OS    RDBKAIWK                         ·
· _ 3      MENSA37          7      OS    RDBKESI3                         ·
···························································································
```

Figure 15-6 Add PCIe Function: Define Access List - ISM

7. HCD now shows the `Define Candidate List` panel. In this panel you select which partitions can access the function besides the access partition. In our example, we do *not* select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel where you can see that the function is now defined (Figure 15-7).

```
------------------------------------------------------------------------------
                              PCIe Function List     Row 1 of 23 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 0040   7F1   _   1   _  ISM           ____  ISM FID=0040 VFID=1 CHID=7F1
```

Figure 15-7 PCIe Function List: Function now created - ISM

8. Now, define the other FID according to the example so far (Figure 15-8).

```
------------------------------------------------------------------------------
                              PCIe Function List     Row 1 of 26 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 0040   7F1   _   1   _  ISM           ____  ISM FID=0040 VFID=1 CHID=7F1
_ 0041   7F1   _   2   _  ISM           ____  ISM FID=0041 VFID=2 CHID=7F1
```

Figure 15-8 PCIe Function List: All Functions now created - ISM

ISM management

From an operational standpoint, SMC-D is like SMC-R. However, SMC-D uses DMA instead of Remote Direct Memory Access (RDMA), and it uses a virtual PCI adapter that is called ISM rather than an RDMA network interface card (RNIC). The ISM interfaces are associated with IP interfaces (for example, HiperSockets (IQD), OSA-Express (OSD) and Network Express Adapters (OSH)), and are dynamically created, automatically started and stopped, and auto-discovered.

SMC-D over ISM does not use queue pair (QP) technology like SMC-R. Therefore, links and Link Groups based on QPs (or other hardware constructs) are not applicable to ISM. The SMC-D protocol has a design concept of a “logical point-to-point connection” that is called an SMC-D link.

> **Note:** The SMC-D information in the **`netstat`** command output is related to ISM link information (not Link Groups).

15.2.3 Defining a NETH PCIe function

As described earlier in this chapter, RoCE adapters along with PCIe functions ROCE and ROC2 are no longer supported on IBM z17. This section provides information about configuring the new NETH PCIe function to enable SMC-R in the Network Express Adapters of the IBM z17.

NETH Overview

An OSH channel path and NETH functions can be defined on the same port of a Network Express adapter. Each port is a separate PCHID and for this reason functions of type NETH does not support the PORT parameter. By allowing OSH and NETH to co-exist on the same port, z/OS Communications Server can initialize a SMC-R network using a single port. In this chapter we cover the NETH definitions. OSH definitions are covered in 13.4, “Network Express for channel type OSH” on page 287.

On IBM z17 there is a maximum of 4096 NETH PCIe functions per CPC and up to 123 NETH virtual functions (VF) can be defined per PCHID.

NETH Configuration

In the following example we define two NETH PCIe functions in two Network Express Adapter ports (each port is a different PCHID) for LPAR MENSA32 using PNETID REDBOOKS129. The function definitions are:

- `PCHID=1B0` to FID 3432 (`VF=1`) for LPAR MENSA32 on `CPC=MENSA`
- `PCHID=278` to FID 3632 (`VF=1`) for LPAR MENSA32 on `CPC=MENSA`

Complete the following steps:

1. From the main HCD panel, select option `1.3. Processor List`.
2. Enter `f` (work with PCIe functions) next to the processor (`MENSA`) to which you want to define the NETH functions, and press Enter (see Figure 15-9).

```
-------------------------------------------------------------------------------
                                Processor List       Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
f MENSA    9175     ME1      LPAR  0B9FB89175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
******************************* Bottom of data ********************************
```

Figure 15-9 Processor List: Work with PCIe functions - NETH

3. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (see Figure 15-10).

```
------------------------------------------------------------------------------
                              PCIe Function List     Row 1 of 22 More:       >
Command ===> add____________________________________________ Scroll ===> PAGE

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
```

Figure 15-10 PCIe Function List: Add line command - NETH

4. Make the following updates (see Figure 15-11 on page 340), and press Enter:
   - Update `Function ID` to `3432`.
   - Update `Type` to `NETH`.
   - Update `Channel ID` to `1B0`.
   - Update `Virtual Function ID` to `1`.
   - Update `Description` to the description that you want.

```
............................. Add PCIe Function .............................
.                                                                            .
.                                                                            .
. Specify or revise the following values.                                    .
.                                                                            .
. Processor ID . . . . . . . : MENSA        Mensa                            .
.                                                                            .
. Function ID  . . . . . . . . 3432                                          .
. Type . . . . . . . . . . . . NETH________  +                               .
.                                                                            .
. Channel ID . . . . . . . . . 1B0  +                                        .
. Port . . . . . . . . . . . . _    +                                        .
. Virtual Function ID  . . . . 1__  +     Physical Function    _  (Y/N)      .
. Number of virtual functions  1                                             .
. UID  . . . . . . . . . . . . ____                                          .
.                                                                            .
. Description  . . . . . . . . NETH FID=3432 VFID=1 CHID=1B0___              .
.                                                                            .
.                                                                            .
.                                                                            .
..............................................................................
```

Figure 15-11 PCIe Function List: Add PCIe function - NETH

5. Answer NO in the Allow Promiscuous Mode panel (see Figure 15-12)

```
···················· Allow Promiscuous Mode ···················
·                                                             ·
·                                                             ·
· Specify Yes to indicate that the FID is allowed to use      ·
· promiscuous mode when enabled by the operating system,      ·
· otherwise specify No.                                       ·
·                                                             ·
· Promiscuous Mode . . . . . . . . . . NO                     ·
·                                                             ·
·                                                             ·
·                                                             ·
···············································································
```

Figure 15-12 Add PCIe Function: Allow Promiscuous Mode - NETH

Note: The new OSH CHIPD of IBM z17 uses EQDIO instead of the legacy QDIO. The option to allow a function to be in promiscuous mode is necessary to support existing client capability, since Linux does not support EQDIO.

6. Update `Physical network ID` 1 to `REDBOOKS129` in the Add/Modify Physical Network IDs panel (see Figure 15-13), and press Enter.

```
····················· Add/Modify Physical Network IDs ·····················
·                                                                          ·
·                                                                          ·
· If the Channel ID (CHID) is associated to one or more physical networks, ·
· specify each physical network ID corresponding to each applicable        ·
· physical port.                                                           ·
·                                                                          ·
· Physical network ID 1  . . REDBOOKS129_____                              ·
· Physical network ID 2  . . ________________                              ·
· Physical network ID 3  . . ________________                              ·
· Physical network ID 4  . . ________________                              ·
·                                                                          ·
·                                                                          ·
·                                                                          ·
·                                                                          ·
·····························································································
```

Figure 15-13 Add PCIe Function: Add/Modify Physical Network IDs - NETH

7. Select the required Access LPAR for Function access list. In our example, we use LPAR MENSA32(OS). Press Enter (see Figure 15-14).

```
··························· Define Access List ···························
·                                                             Row 29 of 57 ·
· Command ===> _________________________________________ Scroll ===> CSR    ·
·                                                                          ·
· Select one partition for the access list.                                ·
·                                                                          ·
· Function ID  . . . . : 3432                                              ·
·                                                                          ·
· / CSS ID Partition Name   Number Usage Description                       ·
· _ 2      MENSA29          9      OS                                      ·
· _ 3      MENSA3A          A      OS    WTSCMXA                           ·
· _ 3      MENSA3B          B      OS    RDBKPOK                           ·
· _ 3      MENSA3C          C      CF    Redbooks CF86                     ·
· _ 3      MENSA3D          D      CF    Redbooks CF87                     ·
· _ 3      MENSA3E          E      CF    Redbooks CF76                     ·
· _ 3      MENSA3F          F      CF    Redbooks CF77                     ·
· _ 3      MENSA31          1      OS    RDBKCRYP                          ·
· / 3      MENSA32          2      OS    WTSC75                            ·
· _ 3      MENSA33          3      OS    RDBKAIDV                          ·
· _ 3      MENSA34          4      OS    WTSC80                            ·
· _ 3      MENSA35          5      OS    WTSC81                            ·
· _ 3      MENSA36          6      OS    RDBKAIWK                          ·
· _ 3      MENSA37          7      OS    RDBKESI3                          ·
····························································
```

Figure 15-14 Add PCIe Function: Define Access List - NETH

8. HCD now shows the `Define Candidate List` panel. In this panel you select which partitions can access the function besides the access partition. In our example, we do *not* select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel where you can see the function that is defined (see Figure 15-15).

```
------------------------------------------------------------------------------
                              PCIe Function List    Row 25 of 25 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 3432   1B0   _    1  _  NETH          ____  NETH FID=3432 VFID=1 CHID=1B0
****************************** Bottom of data *******************************
```

Figure 15-15 PCIe Function List: Function now created - NETH

9. Define the other FID according to the example so far (see Figure 15-16).

```
------------------------------------------------------------------------------
                              PCIe Function List    Row 25 of 26 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA       Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 3432   1B0   _   1   _  NETH          ____  NETH FID=3432 VFID=1 CHID=1B0
_ 3632   278   _   1   _  NETH          ____  NETH FID=3632 VFID=1 CHID=278
****************************** Bottom of data ********************************
```

Figure 15-16 PCIe Function List: All Functions now created - NETH

NETH Management

This section introduces the z/OS commands that are related to the NETH PCIe functions, and shows the responses on our test system.

DISPLAY PCIE command

You can use the **DISPLAY PCIE** command to display these items:

- All registered device drivers (with assigned printable names).
- All available or in-use PCIe functions and their associated device types.
- Information about a specific PCIe device with a list of the client address spaces that use the device.

Example 15-2 is an example of the **DISPLAY PCIE** command. You can confirm the FID and VFN that you defined. The FID is represented as PFIDs. Both FIDs are still in the stand by (STNBY) status. Later we will bring this FIDs online.

Example 15-2 Example of the D PCIE command

```
DISPLAY PCIE
IQP022I 10.30.54 DISPLAY PCIE 072
PCIE     0012 ACTIVE
PFID      DEVICE TYPE NAME          STATUS  ASID  JOBNAME  CHID VFN  PN
00002332  8GB zHyperLink            ALLC    0019  IOSAS    0218 0003 1
00002432  8GB zHyperLink            ALLC    0019  IOSAS    0218 0003 2
00003632  Network Express           STNBY                  0278 0001
00003432  Network Express           STNBY                  01B0 0001
00002132  8GB zHyperLink            ALLC    0019  IOSAS    01FC 0003 1
00002232  8GB zHyperLink            ALLC    0019  IOSAS    01FC 0003 2
```

Example 15-3 is an example of the **DISPLAY PCIE,PFID=pfid** command. After you define the new PCIe function, you can see more details for this particular FID with this command.

Example 15-3 Example of the DISPLAY PCIE,PFID=pfid command

```
DISPLAY PCIE,PFID=3632
IQP024I 10.34.56 DISPLAY PCIE 074
PCIE     0012 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
00003632  Network Express          STNBY                  0278 0001
 CLIENT ASIDS: NONE
 PNetID 1: REDBOOKS129
```

Example 15-4 is an example of the **DISPLAY PCIE,DD** command. You can confirm the details of the device drivers that are installed in the system. In the display output we can see the new Network Express driver.

Example 15-4 Example of the DISPLAY PCIE,DD command

```
DISPLAY PCIE,DD
IQP023I 10.41.13 DISPLAY PCIE 088
PCIE     0012 ACTIVE
DEV TYPE   DEVICE TYPE NAME
10140613   8GB zHyperLink
15B36750   10GbE RoCE
15B31003   10GbE RoCE
15B31004   10GbE RoCE Express
15B31016   RoCE Express2
15B31014   RoCE Express2
15B3101E   Network Express
101404ED   ISM
```

The CONFIG command

You can use the **CONFIG** command to bring the PFID online or offline.

Example 15-5 is an example of a **CONFIG PFID(xx),ONLINE** command.

Example 15-5 Example of a CONFIG PFID(xx),ONLINE command

```
CONFIG PFID(3632),ONLINE
IEE504I PFID(3632),ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Example 15-6 is an example of a **CONFIG PFID(xx),OFFLINE** command.

Example 15-6 Example of a CF PFID(xx),OFFLINE command

```
CONFIG PFID(3632),OFFLINE
IEE505I PFID(3632),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

15.2.4 Defining a zHyperLink PCIe function

The zHyperLink Express is a direct-connect, short-distance IBM Z I/O feature that works with a High-Performance FICON storage area network (SAN) infrastructure.

IBM zHyperLink reduces latency by interconnecting the IBM z17 directly to the I/O bay of the DS8880 storage system or later.

zHyperLink Overview

zHyperLink uses a PCIe feature that is called zHyperLink Express2.0 (Feature Code #0351).

There are two ports per feature, and up to 127 VFIDs can be defined per port.

zHyperLink configuration

In this example, we define these items:

- `CHID=1FC` to `FID = 2132, VFIDs = 3, Port = 1,` on `CPC = MENSA`
- `CHID=1FC` to `FID = 2232, VFIDs = 3, Port = 2,` on `CPC = MENSA`
- `CHID=218` to `FID = 2332, VFIDs = 3, Port = 1,` on `CPC = MENSA`
- `CHID=218` to `FID = 2432, VFIDs = 3, Port = 2,` on `CPC = MENSA`

Complete the following steps:

1. From the main HCD panel, select option `1.3. Processor List.` Enter `f` (work with PCIe functions) next to the processor (`MENSA`) to which you want to define the zHyperLink functions, and press Enter (see Figure 15-17).

```
-------------------------------------------------------------------------------
                                Processor List       Row 1 of 3 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
f MENSA    9175     ME1      LPAR  0B9FB89175 Mensa
_ PAVO     3931     A01      LPAR  071A083931 Pavo
_ VELA     3932     AGZ      LPAR  087F283932 Vela
******************************* Bottom of data ********************************
```

Figure 15-17 Processor List: Work with PCIe functions - zHyperLink

2. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (Figure 15-18).

```
------------------------------------------------------------------------------
                              PCIe Function List     Row 1 of 22 More:       >
Command ===> add____________________________________________ Scroll ===> CSR

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA        Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
```

Figure 15-18 PCIe Function List: Add line command - zHyperLink

3. Make the following updates (Figure 15-19), and press Enter:
   - Update `Function ID` to `2132`.
   - Update `Type` to `ZHYPERLINK`.
   - Update `Channel ID` to `1FC`.
   - Update `Port ID` to `1`.
   - Update `Virtual Function ID` to 3.
   - Update `Description` to the description that you want.

```
····························· Add PCIe Function ·····························
·                                                                             ·
·                                                                             ·
· Specify or revise the following values.                                     ·
·                                                                             ·
· Processor ID . . . . . . . : MENSA        Mensa                             ·
·                                                                             ·
· Function ID  . . . . . . . . 2132                                           ·
· Type . . . . . . . . . . . . ZHYPERLINK    +                                ·
·                                                                             ·
· Channel ID . . . . . . . . . 1FC  +                                         ·
· Port . . . . . . . . . . . . 1    +                                         ·
· Virtual Function ID  . . . . 3    +     Physical Function    _  (Y/N)       ·
· Number of virtual functions  1                                              ·
· UID  . . . . . . . . . . . . ____                                           ·
·                                                                             ·
· Description  . . . . . . . . ZHYL FID=2132 VFID=3 CHID=1FC                  ·
·                                                                             ·
·                                                                             ·
·                                                                             ·
·······························································································
```

Figure 15-19 PCIe Function List: Add PCIe Function - zHyperLink

4. Select the required `Access LPAR for Function` access list. In our example, we use LPAR `MENSA32(OS)`. Press Enter (see Figure 15-20).

```
··························· Define Access List ···························
·                                                              Row 29 of 57 ·
· Command ===> _________________________________________ Scroll ===> CSR    ·
·                                                                           ·
· Select one partition for the access list.                                 ·
·                                                                           ·
· Function ID  . . . . : 2132                                               ·
·                                                                           ·
· / CSS ID Partition Name   Number Usage Description                        ·
· _ 2      MENSA29          9      OS                                       ·
· _ 3      MENSA3A          A      OS    WTSCMXA                            ·
· _ 3      MENSA3B          B      OS    RDBKPOK                            ·
· _ 3      MENSA3C          C      CF    Redbooks CF86                      ·
· _ 3      MENSA3D          D      CF    Redbooks CF87                      ·
· _ 3      MENSA3E          E      CF    Redbooks CF76                      ·
· _ 3      MENSA3F          F      CF    Redbooks CF77                      ·
· _ 3      MENSA31          1      OS    RDBKCRYP                           ·
· / 3      MENSA32          2      OS    WTSC75                             ·
· _ 3      MENSA33          3      OS    RDBKAIDV                           ·
· _ 3      MENSA34          4      OS    WTSC80                             ·
· _ 3      MENSA35          5      OS    WTSC81                             ·
· _ 3      MENSA36          6      OS    RDBKAIWK                           ·
· _ 3      MENSA37          7      OS    RDBKESI3                           ·
·············································································
```

Figure 15-20 Add PCIe Function: Define Access List - zHyperLink

5. HCD now shows the `Define Candidate List` panel. In this panel you select which partitions can access the function besides the access partition. In our example, we do *not* select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel, where you can see the function now defined (see Figure 15-21).

```
-------------------------------------------------------------------------------
                              PCIe Function List     Row 7 of 23 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA       Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 2132   1FC   1   3   _  ZHYPERLINK    ____  ZHYL FID=2132 VFID=3 CHID=1FC
```

Figure 15-21 PCIe Function List: Function now created - zHyperLink

6. Now, define the other FIDs according to the example so far (see Figure 15-22).

```
------------------------------------------------------------------------------
                              PCIe Function List    Filter Mode. More:       >
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more PCIe functions, then press Enter. To add, use F11.

Processor ID . . . . : MENSA       Mensa



/ FID    CHID+ P+  VF+ PF Type+         UID   Description
_ 2132   1FC   1   3   _  ZHYPERLINK    ____  ZHYL FID=2132 VFID=3 CHID=1FC
_ 2232   1FC   2   3   _  ZHYPERLINK    ____  ZHYL FID=2232 VFID=3 CHID=1FC
_ 2332   218   1   3   _  ZHYPERLINK    ____  ZHYL FID=2332 VFID=3 CHID=218
_ 2432   218   2   3   _  ZHYPERLINK    ____  ZHYL FID=2432 VFID=3 CHID=218
****************************** Bottom of data ********************************
```

Figure 15-22 PCIe Function List: All Functions now created - zHyperLink

Managing zHyperLink Express

At the time of the writing of this book, the only two exploiters of zHyperLinks are IBM Db2 and IBM MQ. Db2 can use zHyperLinks for both database reads and active logs writes. MQ can only uses zHyperLinks for active log writing. Db2 12 is the first version of Db2 o support zHyperLinks. MQ started supporting zHyperLink in release 9.4. In order to exploit zHyperLink in either Db2 or MQ you must first enable it in z/OS.

z/OS

To enable z/OS for zHyperLink, the **`ZHPF=YES`** and **`ZHYPERLINK OPER=ALL`** statements must be added to the `IECIOSxx` parmlib member, as shown in Example 15-7. In z/OS 3.1 the zHyperlink facility is disabled on the system by default.

Example 15-7 IECIOSxx parmlib enabled for zHyperLink

```
BROWSE    SYS1.PARMLIB(IECIOS00) - 01.20          Line 0000000000 Col 001 080
 Command ===>                                                  Scroll ===> CSR
********************************* Top of Data **********************************
HYPERPAV=XPAV
ZHPF=YES
ZHYPERLINK,OPER=ALL
******************************** Bottom of Data ********************************
```

Note: z/OS 3.1 introduces a new parameter **READLIMIT=n** in the ZHYPERLINK statement of IECIOSxx parmlib member. zHyperLink read is enabled for Db2 objects with CI size less than or equal to the current READLIMIT setting in KB. If no value is specified, like in the above example, then the default value is used. This default depends on the storage controller attached to the channel. For the original zHyperLink capability, the default is 4K (+32 bytes). For DS8K G10 and future offering, the default is 16K (+32 bytes). For more information about the READLIMIT parameter see *MVS Initialization and Tuning Reference* for z/OS 3.1

This process can also be done dynamically by entering the **`SETIOS ZHYPERLINK,OPER=ALL`** console command. The corresponding display command is **`DISPLAY IOS,ZHYPERLINK`**, as shown in Example 15-8.

Example 15-8 DISPLAY IOS,ZHYPERLINK

```
D IOS,ZHYPERLINK
IOS634I 11.25.29 IOS SYSTEM OPTION 311
ZHYPERLINK IS ENABLED FOR READ AND WRITE OPERATIONS
ZHYPERLINK READ DATA LIMIT IS SET TO THE DEFAULT
```

The **DISPLAY PCIE** command can be used to display the available PFIDs for zHyperLink, as shown in Example 15-9.

Example 15-9 DISPLAY PCIE

```
DISPLAY PCIE
IQP022I 10.30.54 DISPLAY PCIE 072
PCIE     0012 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
00002332  8GB zHyperLink           ALLC    0019  IOSAS    0218 0003 1
00002432  8GB zHyperLink           ALLC    0019  IOSAS    0218 0003 2
00003632  Network Express          STNBY                  0278 0001
00003432  Network Express          STNBY                  01B0 0001
00002132  8GB zHyperLink           ALLC    0019  IOSAS    01FC 0003 1
00002232  8GB zHyperLink           ALLC    0019  IOSAS    01FC 0003 2
```

Example 15-10 shows the **DISPLAY PCIE,PFID=pfid** command to display a specific zHyperLink PFID.

Example 15-10 DISPLAY PCIE,PFID=pfid

```
D PCIE,PFID=2132
IQP024I 12.00.22 DISPLAY PCIE 355
PCIE     0012 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
00002132  8GB zHyperLink           ALLC    0019  IOSAS    01FC 0003 1
 CLIENT ASIDS: NONE
 CU WWNN: 500507630BFFD055  CU Link Id: 0081
 S/W State: Allocated
 Port State: Operational
 CU Node Descriptor: 002107.998.IBM.75.0000000LTV61
```

The results of running the **DISPLAY M=CU(cun)** command against a CU that is enabled for zHyperLink are shown in Example 15-11.

Example 15-11 DISPLAY M=CU(cun)

```
D M=CU(696)
IEE174I 09.35.22 DISPLAY M 859
CONTROL UNIT 0696
CHP                   40   41   42   43
ENTRY LINK ADDRESS    0112 0212 0115 0215
DEST LINK ADDRESS     0103 0203 0104 0204
CHP PHYSICALLY ONLINE Y    Y    N    Y
PATH VALIDATED        Y    Y    N    Y
MANAGED               N    N    N    N
ZHPF - CHPID          Y    Y    Y    Y
ZHPF - CU INTERFACE   Y    Y    N    Y
INTERFACE ID          0100 0030 .... 0230
CONNECTION SECURITY   None None None None
```

```
MAXIMUM MANAGED CHPID(S) ALLOWED = 0
DESTINATION CU LOGICAL ADDRESS = 96
CU ND            = 002107.998.IBM.75.0000000LTV61.0100
CU NED           = 002107.998.IBM.75.0000000LTV61.9600
TOKEN NED        = 002107.900.IBM.75.0000000LTV61.9600
WWNN             = 500507630BFFD055
FUNCTIONS ENABLED = ZHPF, ZHYPERLINK, XPAV
XPAV CU PEERS    = 0692, 0694, 0696, 0698
DEFINED DEVICES
  08400-084EF
DEFINED PAV ALIASES
  084F0-084FF
USABLE HYPERPAV ALIASES = 16
ZHYPERLINK READ DATA LIMIT IS 4K
ZHYPERLINKS
  PFID      PCHID  Port  LinkId  S/W St   Port St
  00002132  01FC   01    0081    Alloc    Oper
  00002232  01FC   02    0181    Alloc    Oper
  00002332  0218   01    0281    Alloc    Oper
  00002432  0218   02    0381    Alloc    Oper
```

The results for the **DISPLAY M=DEV(devno)** command against a device that is enabled for zHyperLink are shown in Example 15-12.

Example 15-12 DISPLAY M=DEV(devno)

```
D M=DEV(8400)
IEE174I 09.46.40 DISPLAY M 863
DEVICE 08400   STATUS=ONLINE
CHP                   40   41   42   43
ENTRY LINK ADDRESS    0112 0212 0115 0215
DEST LINK ADDRESS     0103 0203 0104 0204
PATH ONLINE           Y    Y    N    Y
CHP PHYSICALLY ONLINE Y    Y    N    Y
PATH OPERATIONAL      Y    Y    N    Y
MANAGED               N    N    N    N
CU NUMBER             0696 0696 0696 0696
INTERFACE ID          0100 0030 .... 0230
CONNECTION SECURITY   None None None None
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 96
SCP CU ND        = 002107.998.IBM.75.0000000LTV61.0100
SCP TOKEN NED    = 002107.900.IBM.75.0000000LTV61.9600
SCP DEVICE NED   = 002107.900.IBM.75.0000000LTV61.9600
WWNN             = 500507630BFFD055
HYPERPAV ALIASES CONFIGURED = 16
ZHYPERLINKS AVAILABLE = 4
FUNCTIONS ENABLED = MIDAW, ZHPF, XPAV, ZHYPERLINK
```

This display command is enhanced with the new parameter **ZHYPERLINK** to show whether the device can use zHyperLink. The response is shown for a device that can use zHyperLink in an z/OS environment that is enabled for read/write, as shown in Example 15-13.

Example 15-13 DISPLAY M=DEV(devno),ZHYPERLINK: Device enabled for zHyperLink

```
D M=DEV(8400),ZHYPERLINK
```

```
IEE587I 09.48.48 DISPLAY M 869
DEVICE 08400   STATUS=ONLINE
DEVICE IS ENABLED FOR ZHYPERLINK
READ AND WRITE OPERATIONS ARE ENABLED
ZHYPERLINKS AVAILABLE = 4
ZHYPERLINK READ DATA LIMIT IS 4K
```

The result for a device that cannot use zHyperLink is shown in Example 15-14.

Example 15-14 DISPLAY M=DEV(devno),ZHYPERLINK: Device not enabled for zHyperLink

```
DISPLAY M=DEV(8400),ZHYPERLINK
IEE587I 09.48.30 DISPLAY M 867
DEVICE 08000   STATUS=OFFLINE
DEVICE IS DISABLED FOR ZHYPERLINK FOR THE FOLLOWING REASON(S):
  THERE ARE NO ZHYPERLINKS AVAILABLE
```

Db2

To enable IBM Db2 to use zHyperLinks, the Db2 zParm must be modified to enable zHyperLink, as shown in the example in Figure 15-23.

```
DSNIIPA3             INSTALL Db2 - DATA PARAMETERS PANEL
===>

Check parameters and reenter to change:
 1 PERMANENT UNIT NAME    ==> 3390     Device type for MVS catalog and
                                       partitioned data sets
 2 TEMPORARY UNIT NAME    ==> SYSDA    Device type for temporary data sets
 3 Db2 zHyperlinks SCOPE  ==> ENABLE   Scope of zHyperlinks I/O connections:
                                       (ENABLE, DISABLE, DATABASE, LOG)

                                         ---------------- SMS -----------------
                              VOL/SER    DATA CLASS    MGMT CLASS    STOR CLASS
                              -------    ----------    ----------    ----------
 4 CLIST ALLOCATION       ==> ______  ==> ________  ==> ________  ==> ________
 5 NON-VSAM DATA          ==> ______  ==> ________  ==> ________  ==> ________
 6 VSAM CATALOG, DEFAULT, ==> ______  ==> ________  ==> ________  ==> ________
   AND WORK FILE DATABASE
 7 LOG COPY 1, BSDS 2     ==> ______  ==> ________  ==> ________  ==> ________
 8 LOG COPY 2, BSDS 1     ==> ______  ==> ________  ==> ________  ==> ________



PRESS:  ENTER to continue   RETURN to exit   HELP for more information
```

Figure 15-23 Db2 zParm: zHyperLink

The acceptable values for the Db2 13 zHyperLink Scope are the following ones:

`ENABLE` Db2 requests the zHyperLink protocol for all eligible I/O requests.

`DISABLE` Db2 does not use the zHyperLink for any I/O requests.

`DATABASE`	Db2 requests the zHyperLink protocol for only database synchronous read I/Os.
`ACTIVELOG`	Db2 requests the zHyperLink protocol only for active log-write I/Os.

Note: In order for the zHyperLink protocol to be enabled for I/O requests by the Db2 system, in addition to specifying a value of ENABLE, DATABASE, or ACTIVELOG for the ZHYPERLINK parameter, the SMS storage class of the Db2 data sets must be enabled for zHyperLink eligibility using the Integrated Storage Management Facility (ISMF).

For more information about the usage of zHyperLink in Db2 see *Installing or Migrating to Db2 13*.

MQ

You can configure IBM MQ to use zHyperLink by using one of the following methods:

- Specify ZHYLINK(YES) in the log parameters.
- Issue the command SET LOG ZHYLINK(YES).

For more information about MQ use of zHyperlink see *Configuring Queue Managers on z/OS*.

16

Defining Channel to Channel Connections

This chapter describes the configuration steps for fibre connection (FICON) channel-to-channel connections (CTC). CTCs can be needed for VTAM or GRS communications or for z/VM Single System Image communications.

Naming: This publication covers the IBM z17 ME1. Throughout this chapter, we might refer to these machines as IBM z17.

This chapter includes the following topics:

- Understanding FCTC
- Defining FCTC
- Verifying FCTC

16.1 Understanding FCTC

FCTC is used to connect two logical partitions directly using FICON. The connected partitions can reside on the same central processor complex (CPC) or on different CPCs. FCTC is used for communication between these partitions. Possible exploiters are for example VTAM and GRS in z/OS or the single system image (SSI) cluster of z/VM.

Both sides of an FCTC connection use a Fibre Channel (FC) channel path that is defined to one or more FCTC CUs representing the target systems (logical partitions) for the CTC connections.

A logical address (CU address or CUADD) is required in the definition of the control unit representing the target system when the FC channel path is attached to a shared FC channel path to identify the LPAR to communicate with. This CUADD has to be the number of the target LPAR prefixed with the channel subsystem (CSS) number of the target LPAR.

When attached to an unshared FC channel path, the CUADD must be zero or not specified. For more background information, see the CTC Definitions page.

FCTC communication within a single CPC can be accomplished with a single FC channel path connecting to an FC switch. Both sides of the CTC connection are represented by the same single FC channel path. This configuration results in the following advantages:

- Reduced number of channels required
- Simplified configuration design
- Simplified configuration definition

The following considerations apply to all FCTC configurations:

- The server at each end of an FCTC connection uses a FICON native (CHPID type FC) channel.
- The FICON native channel at each end of the CTC connection has an FCTC control unit (CU) defined.
- The FCTC devices on the FCTC CU are defined as type FCTC.
- The FICON native channel at each end of the FCTC connection supports the FCTC CUs.

In an FCTC configuration, FCTC CUs are defined at each end, but only one end will provide the FCTC CU function during operation. During initialization of the logical connection between two ends of an FCTC connection, the channel that provides the FCTC CU function is determined by using an algorithm.

For more information about how to implement FCTC, see the following publications:

- *FICON Planning and Implementation Guide*, SG24-6497
- *FICON CTC Implementation,* REDP-0158
- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804

16.2 Defining FCTC

In HCD select your work IODF and go to 1.3 (Processors) and there select the CPC with "S" and then the CSS containing the FCTC channel also with "S" to see the Channel Path List. Add (F11) the needed CHPID if it is not yet defined or change ("C”) the definition of the

CHPID that is to be used as FCTC like shown in Figure 16-1 on page 355. It has to be a FICON CHPID (FC).

```
+-------------------- Change Channel Path Definition ---------------------+
|                                                                         |
|                                                                         |
| Specify or revise the following values.                                 |
|                                                                         |
| Processor ID . . . . : MENSA       Mensa                                |
| Configuration mode . : LPAR                                             |
| Channel Subsystem ID : 0                                                |
|                                                                         |
| Channel path ID  . . . . 44    +          Channel ID  12C  +            |
| Channel path type  . . . FC    +                                        |
| Operation mode . . . . . SPAN  +                                        |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +     |
| Description  . . . . . . FCTC 32G LX                                    |
|                                                                         |
| Specify the following values only if connected to a switch:             |
|                                                                         |
| Dynamic entry switch ID  01  + (00 - FF)                                |
| Entry switch ID  . . . . 01  +                                          |
| Entry port . . . . . . . 17  +                                          |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap    |
| F12=Cancel                                                              |
+-------------------------------------------------------------------------+
```

Figure 16-1 Defining the channel for a CTC

After this has been done, select the CHPID with "S" to see the Control Unit List of this CHPID. Here you have to add (F11) a control unit for each target LPAR that shall be connected via this FCTC like shown in Figure 16-2 on page 356. The control unit type has to be FCTC and the CUADD has to be the LPAR number including the CSS number. For example, for LPAR number 11 (hex "B") in CSS number 4 the CUADD would be 4B. You can define up to 256 unit addresses starting with 00 for each control unit.

```
+-------------------------- Add Control Unit ---------------------------+
|                                                                        |
|                                                                        |
| Specify or revise the following values.                                |
|                                                                        |
| Control unit number  . . . . 444B  +                                   |
| Control unit type  . . . . . FCTC_________   +                         |
|                                                                        |
| Serial number  . . . . . . . __________                                |
| Description  . . . . . . . . CTC to MENSA4B__________________          |
|                                                                        |
| Connected to switches  . . . 01  __  __  __  __  __  __  __   +        |
| Ports  . . . . . . . . . . . 17  __  __  __  __  __  __  __   +        |
|                                                                        |
| If connected to a switch:                                              |
|                                                                        |
| Define more than eight ports . . 2   1.  Yes                           |
|                                      2.  No                            |
| Propose CHPID/link addresses and                                       |
| unit addresses . . . . . . . . . 2   1.  Yes                           |
|                                      2.  No                            |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap        |
| F12=Cancel                                                             |
+------------------------------------------------------------------------+
```

Figure 16-2 Defining an FCTC control unit

After defining the control unit you must add device definitions to the control unit. To achieve this, enter "S" in front of the control unit to get to the I/O Device List. Here you have to add

(F11) as many devices of the type FCTC as needed for your purposes. Adding devices is shown in Figure 16-3.

```
+------------------------------- Add Device -------------------------------+
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                   |
|                                                                            |
| Device number  . . . . . . . . 4420  + (0000 - FFFF)                     |
| Number of devices  . . . . . . 2___                                       |
| Device type  . . . . . . . . . FCTC_________  +                           |
|                                                                            |
| Serial number  . . . . . . . . __________                                 |
| Description  . . . . . . . . . CTC to MENSA4B__________________           |
|                                                                            |
| Volume serial number . . . . . ______  (for DASD)                         |
|                                                                            |
| PPRC usage . . . . . . . . . . _  + (for DASD)                           |
|                                                                            |
| Connected to CUs . . 444B  ____  ____  ____  ____  ____  ____  ____  +   |
|                                                                            |
|                                                                            |
|                                                                            |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap     |
| F12=Cancel                                                                 |
+--------------------------------------------------------------------------+
```

Figure 16-3 Defining an FCTC device

16.3 Verifying FCTC

After the IODF with the FCTC configuration has been activated and all physical cabling for the FCTC has been done, there are two places where you can verify that the FCTC connection is operational

- Checking the status using z/OS commands.
- Checking the status using the Service Element (SE)

16.3.1 Checking the status using z/OS commands

If for example the CHPID 44 is defined for FCTC usage and the FTC control unit use devices 4E30 to 4E33, then you can run the z/OS system command **`D M=CHP(44)`**. It should show the CHPID as ONLINE and the devices also as online (indicated by a “+” sign).

16.3.2 Checking the status by using the SE:

- From the HMA, select the CPC (under Systems Management) where the CHPID or PCHID that you want to verify is and click **Single Object Operations** (under the **Recovery** task options). This opens the service element (SE).
- On the SE, select the same CPC and click **Channels**. Look for the PCHID that you are interested in checking the status of. It should be shown as Operational and Online.
- To get more details about the PCHID, click the PCHID to open the details window.

- Another way to check the status is from the LPAR view. Select the LPAR and then the CHPIDs option under that LPAR. You can look for the CHPID and check the status. It should be shown as Operational and Online.
- For more details, click the CHPID.

A

Additional material

This book refers to additional material that can be downloaded from the internet as described in the following sections.

Locating the web materials

The web material that is associated with this book is available in softcopy on the internet from the IBM Redbooks web server:

https://www.redbooks.ibm.com/abstracts/sg248960.html

Alternatively, you can go to the IBM Redbooks website:

ibm.com/redbooks

Search for SG24-8960, select the title, and then click **Additional materials** to open the directory that corresponds with the IBM Redbooks form number, SG24-8960.

Using the web materials

The additional web materials that accompany this book includes the following files:

File name	*Description*
Relocation Services Inventory.pdf	System Discontinuance / Relocation form
8460_DH2_Image_worksheet.xlsx	Worksheet for gathering setup information

Downloading and extracting the web material

Create a subdirectory (folder) on your workstation, and extract the contents of the web material compressed file into this folder.

IBM z17 Configuration Setup

IBM z17 Configuration Setup

IBM z17 Configuration Setup

IBM z17 Configuration Setup

IBM z17
Configuration Setup

IBM z17
Configuration Setup

Abbreviations and acronyms

4HRA	4-hour rolling average
AES	Advanced Encryption Standard
AID	adapter ID
BTS	Backup Time Server
CCA	Common Cryptographic Architecture
CCN	Configuration Control Number
CE LR	Coupling Express2 LR
CF	Coupling Facility
CFACT	Coupling Facility Activity
CFCC	Coupling Facility Control Code
CFOVER	Coupling Facility Overview
CHID	channel ID
CHPID	channel path ID
CMPCS	compression co-processor
CMT	CHPID Mapping Tool
CP	central processor
CPACF	CP Assist for Cryptographic Functions
CPC	central processor complex
CSS	channel subsystem
CST	Coordinated Server Time
CTC	channel-to-channel
CTN	Coordinated Timing Network
CTS	Current Time Server
CU	control unit
CUADD	control unit address
DASD	direct access storage device
DES	Data Encryption Standard
DFT	distributed function terminal
DHD	defer host disconnect
DLA	Destination Link Address
DMA	Direct Memory Access
DPM	Dynamic Partition Manager
DRNG	Deterministic Random Number Generation
DREQ	Direct Memory Access request
DWDM	Dense Wavelength Division Multiplexing
EDT	Eligible Device Table
ETS	External Time Source
FC	Fibre Channel
FCP	Fibre Channel Protocol
FCTC	Fibre Connection channel-to-channel
FICON	Fibre Connection
FID	function ID
FINRA	Financial Industry Regulatory Authority
GbE	Gigabit Ethernet
Gbps	gigabytes per second
GRS	global resource serialization
HA	high availability
HCD	Hardware Configuration Definition
HCM	Hardware Configuration Manager
HDD	hard disk drive
HMA	Hardware Management Appliance
HMC	Hardware Management Console
IBM	International Business Machines Corporation
IBM SSR	IBM Systems Service Representative
IC	Internal Coupling
ICA	Integrated Coupling Adapter
ICA SR	Integrated Coupling Adapter Short Reach
ICC	Integrated Console Controller
ICF	Internal Coupling Facility
ICSF	Integrated Cryptographic Service Facility
IFP	integrated firmware processor
IML	Initial Microcode Load
IOCDS	input/output configuration data set
IOCP	input/output configuration program
IODF	input/output definition file
IPv4	Internet Protocol Version 4
IQD	Internal Queued Direct
IQDIO	internal queued direct input/output
ISM	Internal Shared Memory
ISPF	Interactive System Productivity Facility
JES2	Job Entry Subsystem 2
KVM	Kernel-based Virtual Machine

LAN	local area network
LCSS	logical channel subsystem
LIC	Licensed Internal Code
LPAR	logical partition
LR	Long Reach
LUN	logical unit number
LX	Long Wave
MAC	Media Access Control
MCB	Message Command Block
MCS	Master Control Service
MES	miscellaneous equipment specification
MIF	Multiple Image Facility
MRB	Message Response Block
NMS	national market system
NTP	Network Time Protocol
OAT	OSA Address Table
OOCoD	On/Off Capacity on Demand
OS	operating system
OSA	Open Systems Adapter
OTC	over-the-counter
PCHID	physical channel ID
PCIe	Peripheral Component Interconnect Express
PDU	power distribution unit
PFC	Priority Flow Control
PFID	PCIe function ID
PKCS	Public Key Cryptography Standards
POR	Power on Reset
PPS	Pulse Per Second
PRNG	Pseudo-Random Number Generation
PTF	program temporary fix
PTP	Precision Time Protocol
PTS	Preferred Time Server
QP	queue pair
RCE	Regional Crypto Enablement
RDMA	Remote Direct Memory Access
RG	Resource Group
RMF	Resource Measurement Facility
RNIC	RDMA network interface card
RoCE	Remote Direct Memory Access over Converged Ethernet
RSI	Relocation Services Inventory
RTO	read timeout
SACF	Stand-alone Coupling Facility
SAN	storage area network
SCM	storage-class memory
SE	Support Element
SMC	Shared Memory Communications
SMC-D	SMC - Direct Memory Access over Internal Shared Memory
SMC-R	SMC - Remote Direct Memory Access over Converged Ethernet Express
SMCAT	Shared Memory Communications Applicability Tool
SNA	Systems Network Architecture
SR	Short Reach
SS	Subchannel Set
SSC	Secure Service Container
STP	Server Time Protocol
SX	Short Wave
TDES	Triple Data Encryption Standard
TKE	Trusted Key Entry
TOD	time of day
TRNG	True Random Number Generation
TSO	Time Sharing Option
UA	Unit Address
VCHID	virtual channel ID
VF	virtual function
VFID	virtual function ID
VFM	Virtual Flash Memory
VLAN	virtual local area network
VM	virtual machine
vPCIe	virtual PCIe
VTAM	Virtual Telecommunications Access Method
WWN	worldwide name
WWPN	worldwide port name
zBNA	Z Batch Network Analyzer
zEDC	zEnterprise Data Compression
zHPF	High Performance FICON for IBM Z
zIIP	IBM Z Integrated Information Processor

SG24-8581-00

ISBN DocISBN