Redbooks
ibm.com/redbooks

IBM z14 Model ZR1 Configuration Setup

Octavian Lascu

Aditya Banuaji

Franco Pinto

Ryotaroh Sawada

Martin Söllig

Bill White

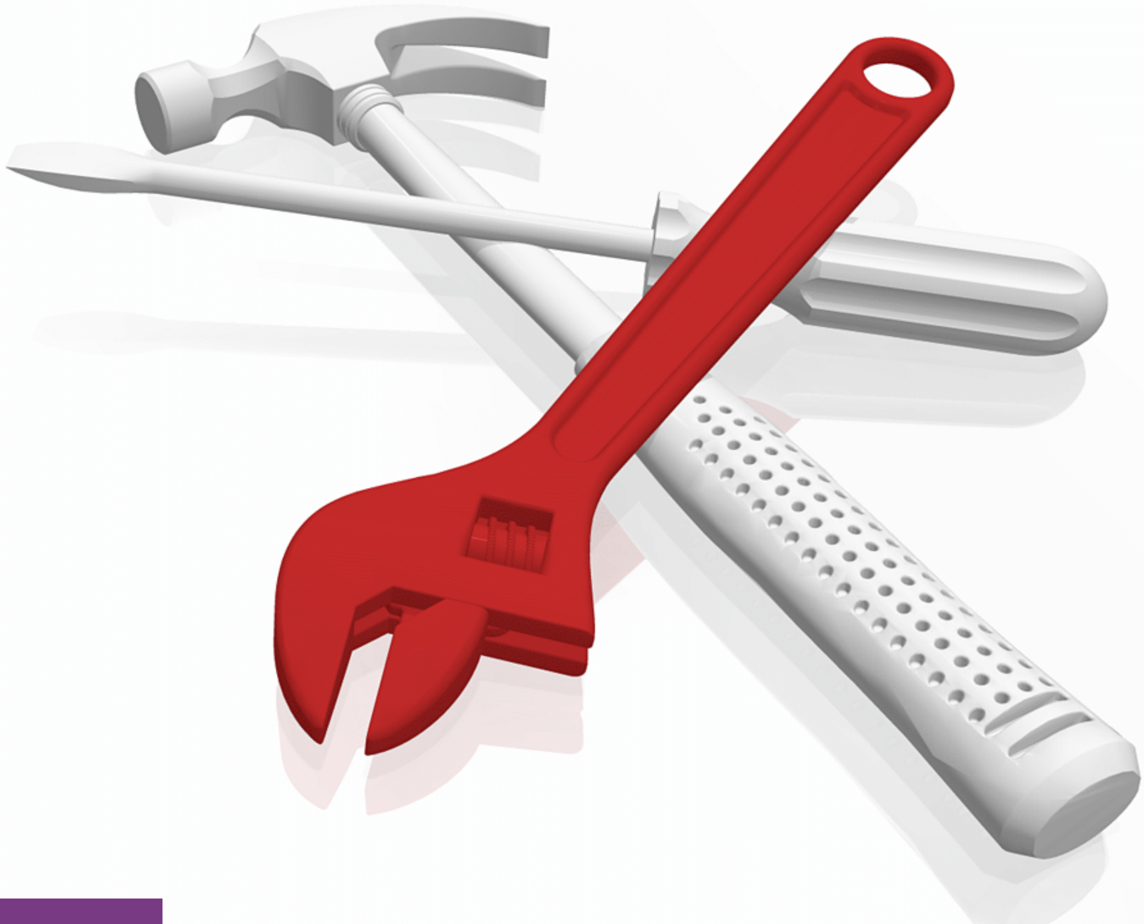

IBM Z

IBM

Redbooks

International Technical Support Organization

IBM z14 Model ZR1 Configuration Setup

November 2019

SG24-8560-00

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

First Edition (November 2019)

This edition applies to IBM z14 Model ZR1 and IBM Z Hardware Management Console Version 2.14.0 (Driver Level 32).

Contents

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 client 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

Preface

This IBM® Redbooks® publication helps you install, configure, and maintain the IBM z14® Model ZR1 (Machine Type 3907). The z14 ZR1 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 generally familiar with current IBM Z technology and terminology. For more information about the functions of the z14 Model ZR1, see *IBM z14 Model ZR1 Technical Introduction*, SG24-8550, and *IBM z14 Model ZR1 Technical Guide*, SG24-8651.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Octavian Lascu is a Senior IT Consultant for IBM Romania with over 25 years of experience. He specializes in designing, implementing, and supporting complex IT infrastructure environments (systems, storage, and networking), including high availability and disaster recovery solutions and high-performance computing deployments. He has developed materials and taught workshops for technical audiences around the world for the past 17 years. He has written several IBM publications.

Aditya Banuaji is a is a Client Technical Sales for IBM Z in Indonesia. His expertise includes IBM Z hardware capacity planning, z/OS®, and Linux on IBM Z. He is involved in building LinuxONE demo environments in Indonesia and leads proof of concept and ISV application porting projects. Aditya holds a master's degree in Electrical Engineering from Seoul National University of Science and Technology, and a bachelor's degree in Electrical Engineering from Bandung Institute of Technology.

Franco Pinto is a Client Technical Specialist in IBM Switzerland. He has over 20 years of experience in the mainframe and IBM z/OS fields. His areas of expertise include IBM Z technical pre-sales covering mainframe sizing and installation planning, and providing support on existing and new IBM Z functions.

Ryotaroh Sawada is an IT Specialist with IBM in Japan. He has 11 years of experience in technical sales and supporting IBM Z clients. His areas of expertise include IBM Z hardware, z/OS and Systems Management on the mainframe. He also co-authored *Extending IBM z/OS System Management Functions with IBM zAware*, SG24-8070.

Martin Söllig is a Consultant IT Specialist in Germany. He has 28 years of experience working in the IBM Z field. He holds a degree in Mathematics from University of Hamburg. His areas of expertise include z/OS and IBM Z hardware, specifically in Parallel Sysplex® and GDPS® environments, and in cryptography on IBM Z.

Bill White is an IBM Redbooks Project Leader and Senior Networking and Connectivity Specialist at IBM Redbooks, Poughkeepsie Center.

Thanks to the authors of the *IBM z14 Configuration Setup*, SG24-8460:

Peter Hoyle
Jannie Houjberg
Kazuhiro Nakajima
Nelson Oliveira
Franco Pinto
Martin Söllig
Bill White

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

Robert Haimowitz
International Technical Support Organization, Poughkeepsie Center

Dale Riedy
Luis Cruz
Barbara Sannerud
IBM Poughkeepsie

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 ITSO 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, International Technical Support Organization
  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 goals of this book. This book is based on scenarios that the team devised based on experience and best practices.

The scenarios implementation is described in this publication, along with the tools that help implement the configurations that are described.

This chapter includes the following topics:

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 an IBM z14 Model ZR1 server. The book covers the planning and preparation tasks that are needed from when a z14 ZR1 is delivered and physically installed up to the point when an LPAR is ready to be activated.

The book is based on two scenarios, and describes the planning considerations and configuration examples in detail from a Hardware Management Console (HMC) and Support Element (SE) and input/output definition file (IODF) perspective.

1.2 Scope

Before you perform the planning and preparation tasks that this book covers, the following tasks must be completed:

- Customer's Configuration Design

  Together with your team, IBM provides design and configuration information for the installation of the z14 ZR1 system that you plan to purchase.

- IBM Order to Manufacturing

  Your IBM representative orders the agreed configuration. IBM makes available for download the machine configuration as a Configuration Report File (CFR). The CFR file can be obtained from the IBM Resource Link website (authentication is required, use your IBM ID) by using a Configuration Control Number (CCN) that is provided by your IBM representative.

- Physical installation

  With support from IBM, the new order or the upgrade to a z14 ZR1 server is physically installed. This process includes the preparations for and activation of the server in Dynamic Partition Mode (DPM).

- HMC/TKE installation

  With support from IBM, the HMCs and the (optional) Trusted Key Entry (TKE) workstations are installed and, if necessary, contents, such as user profiles and API settings, are migrated (in a replacement of HMCs and TKEs).

Based on the flowcharts that are shown in Figures 1-1, Figure 1-2 on page 4, and Figure 1-3 on page 5, this book describes the following scenarios when preparing for a z14 ZR1 server installation:

- Upgrading an existing IBM Z server to a z14 ZR1
- Installing a new z14 ZR1 server
- Installing a new z14 ZR1 server to be managed by using Dynamic Partition Manager

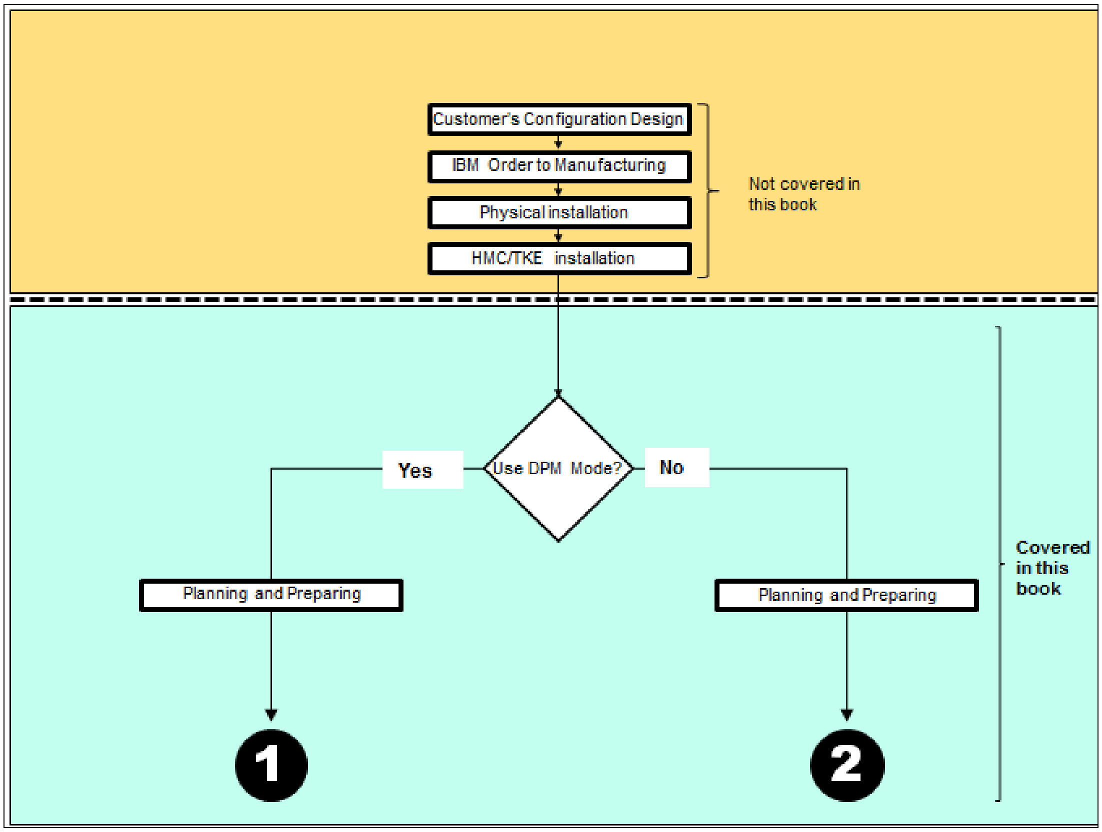

Figure 1-1 Topics that are covered in this book

The flowchart that is shown in Figure 1-1 is divided in two main sets of task streams. The upper stream (marked in light yellow) is based on actions that should be performed before the hardware arrives and are not covered in this book. The lower stream (marked in light green) considers the use of a wizard, such as the configuration method that called Dynamic Partition Mode (DPM), which is described next.

1.3 Dynamic Partition Mode

Linux servers, IBM z/VM® and other specific applications ran on mainframe systems for years, but configuration and setup are fairly complicated and require the use of several specific tools. However, with DPM, system administrators now have a quicker and easier way to deploy Linux servers, z/VM LPARs that host Linux guests, or Secure Service Container (SSC) LEPERs that use only the Hardware Management Console (HMC).

DPM is a configuration manager that is designed for setting up and managing Linux servers that run on a mainframe system. On a DPM-enabled system, the runtime environment for your Linux server, your z/VM Hypervisor, or your SSC application is called a *partition*.

On other platforms, a partition is a portion of the system hard disk drive that you create to run different operating systems on the same disk, or to give the appearance of separate hard disk drives for multiple users or other purposes. On a mainframe system, a partition is a virtual representation of all of the physical hardware resources of that system, which include processors, memory, and input/output (I/O) adapters. On mainframe systems, as on other platforms, an adapter is a physical device that connects the system to other computers or devices.

Note: DPM mode cannot coexist with standard PR/SM mode. When DPM mode is used, only z/VM, Linux, and SSC partitions[a] (LPARs) and connections to them can be defined.

a. For DPM, the term *partition* is equivalent to *Logical Partition (LPAR)*.

The flowchart that is shown in Figure 1-2 is divided in two task streams. The stream on the left side of the flowchart is based on actions that must be performed by IBM on the SEs as preparation before the machine is handed to the customer.

The other stream (right side of the flowchart) is based on definitions that use the DPM wizard in basic mode. Based on the input that is provided to the DPM wizard, a configuration is activated that is used on the z14 ZR1 and the devices that are attached to the server. The actions that are defined in the two streams must be performed in sequence.

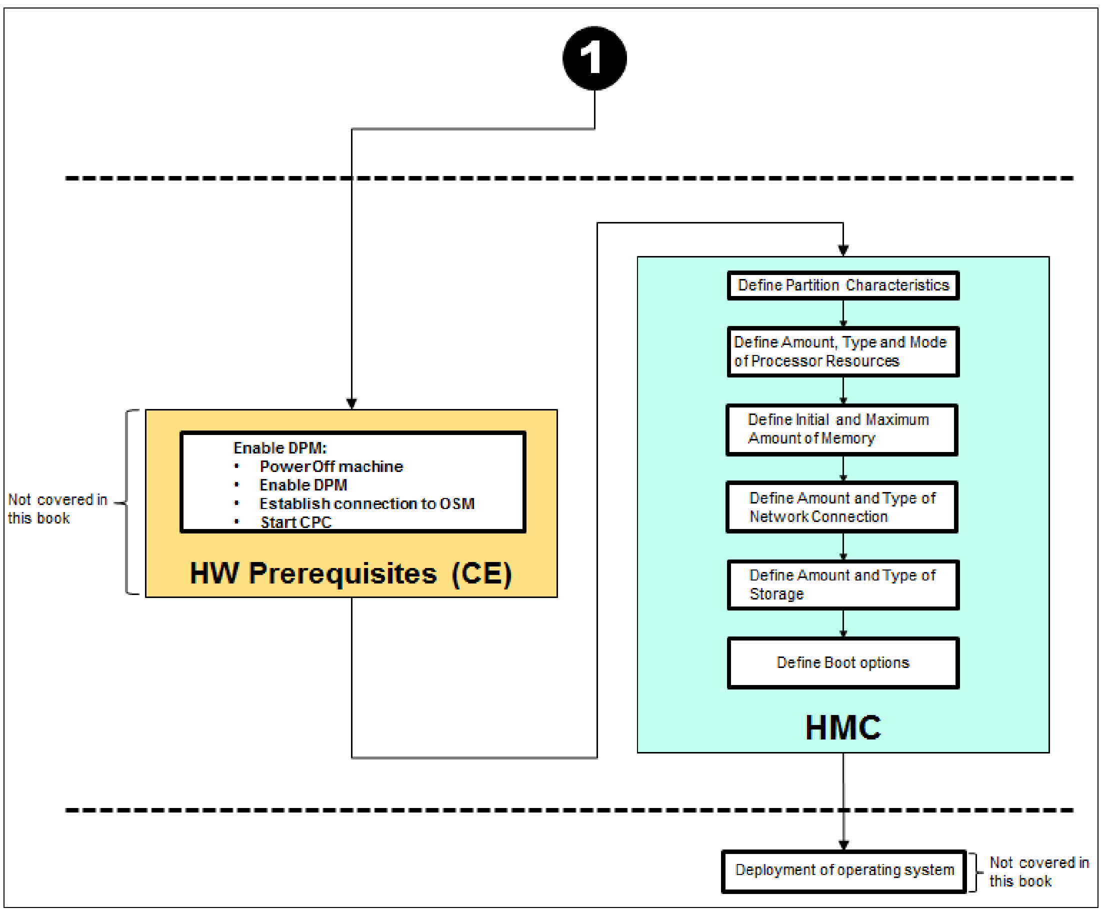

Figure 1-2 Installation flowchart DPM mode (applicable to both scenarios)

The flowchart that is shown in Figure 1-3 is divided in two task streams. One stream (left side of the flowchart) is based on actions that must be performed on the HMC or on the SEs. The other stream (right side of the flowchart) is based on definitions in the DPM wizard.

Based on the DPM wizard, the surrounding I/O configuration data is created that contains a description of all I/O functions and features that are used on the z14 ZR1 and the devices that are attached to the server. The actions that are defined in the two streams can have dependencies in between.

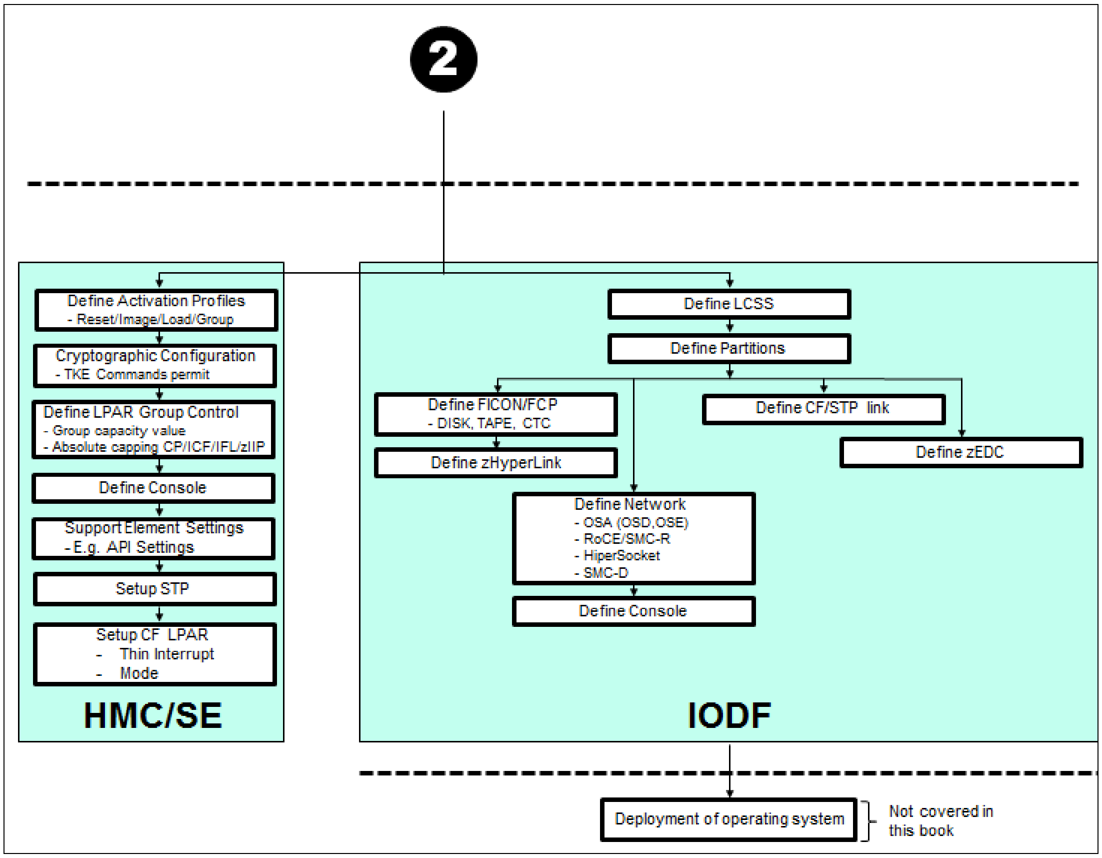

Figure 1-3 Installation flowchart non-DPM mode (applicable to both upgrade/MES and new system)

The IBM Z server is provided with a management appliance (HMC), which is a physical workstation that runs the HMC application. The HMC communicates with the SE appliance (physically installed in the frame of the z14 ZR1), which is used to communicate with the z14 ZR1 hardware. On the HMC/SE, you must set the number of parameters that allows you to activate a number of logical partitions (LPARs) that run an operating system.

To create an IODF (see Figure 1-2 on page 4), a set of tasks must be done 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.

Several HMC/SE and IODF tasks must be planned and prepared. For more information, see *I/O Configuration Using z/OS HCD and HCM*, SG24-7804.

This book describes what is needed from a HMC/SE, IODF, and DPM from the time when the z14 ZR1 is delivered and physically installed up to the point when an LPAR is ready to be activated. The flowcharts are intended to act as a checklist rather than a step-by-step procedure. Nevertheless, 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 operating system (z/OS in this case), see *Mainframe from Scratch: Hardware Configuration and z/OS Build*, SG24-8329.

1.4 Tools

Several tools are provided by IBM that help you to achieve a successful z14 ZR1 server installation. These tools are covered in 2.2, “Tools” on page 9.

Whenever possible, checklists are provided throughout the book that help you go through the steps that are required to complete a specified task.

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

2

Planning considerations for CPC in PR/SM mode

This chapter describes the following scenarios for planning and configuration of an IBM z14 Model ZR1 server by using traditional tools (CPC in PR/SM mode):

- Upgrading or migrating and existing configuration to a z14 ZR1
- Installing a new z14 ZR1

Dynamic Partition Manager: For more information about planning and configuration for a CPC by using Dynamic Partition Manager, see Chapter 16, “Configuring a z14 ZR1 server by using Dynamic Partition Manager” on page 387.

Whenever possible, worksheets that support the planning tasks that are described in this chapter are provided. Throughout this book, we provide various definition examples that use hardware configuration definition (HCD) as the preferred method for configuration. Other tools, such as Hardware Configuration Manager (HCM) and ICP/IOCP, are included for reference only.

This chapter also provides a short overview of tools that IBM provides to help with the complexity of configuring a z14 ZR1, and information about where to obtain the tools and their intended use.

This chapter includes the following topics:

2.1 Scenarios

Throughout this book, we use two distinct scenarios by which we explain the tasks and procedures that are involved to successfully install and configure a z14 ZR1 server.

2.1.1 Scenario 1: Upgrading an existing IBM Z server to a z14 ZR1

This scenario assumes that an IBM Z environment where the IBM Z server is upgraded by using miscellaneous equipment specifications (MES) to a z14 ZR1. The scenario includes a planned outage period for the time of the physical upgrade of the machine. The software environments that are supported by this machine are not available during this period. The serial number of the machine remains the same after the upgrade.

2.1.2 Scenario 2: Installing a new z14 ZR1 server

This scenario assumes that a new z14 ZR1 is installed in a mainframe environment. The z14 ZR1 machine is physically installed along with an existing IBM Z machine. After the installation of the z14 ZR1 is successfully completed and the system is handed over by the IBM service representative, the software environment on the machine to be replaced must be stopped and recabling actions must be performed.

When recabling is complete, postinstallation activities must be performed and the software environment can be brought back online on the new system (z14 ZR1). An outage has still to be planned for this scenario, and a new serial number must be considered, so software keys for the new system must be available.

2.1.3 Differences in planning for the two scenarios

In the first scenario, the physical platform identity to be configured remains the same. No hardware configuration files must be physically migrated to another platform. Because the machine serial number remains the same after the upgrade, no changes to the software licenses are required.

In the second scenario, the physical platform to be configured changes. Hardware configuration files must be prepared on the existing machine, and must be migrated to the new z14 ZR1 server with the attached cabling. The serial number changes with the activation of the z14 ZR1 machine, which means that planning and preparing for software license changes must be considered.

In both scenarios, we assume that bringing up the existing features and functions has highest priority. Adding features and functions that were acquired with the system upgrade or installed in the new z14 ZR1 have a lower priority. The elapsed time of the planned outage can vary significantly, depending on the approach that is chosen in either scenario.

In both scenarios, the following information must be obtained before starting the process of changing to or installing the new z14 ZR1:

- 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 *HCD Users Guide*. SC34-2669.
- The `CFReport` file: The `CFReport` file is downloadable from IBM Resource Link® by entering the Configuration Control Number (CCN). The CCN is provided by your IBM representative.
- The system serial number: If a new z14 ZR1 is installed, a new serial number is provided by your IBM representative.

2.2 Tools

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

The machine types for the current IBM Z platform are listed in Table 2-1.

Table 2-1 Machine types

Server name	**Server short name**	**Machine type (M/T)**
IBM Z z14 Model ZR1	z14 ZR1	3907
IBM Z z14	z14	3906
IBM z Systems® z13s®	z13s	2965
IBM z Systems z13®	z13	2964
IBM zEnterprise® BC12	zBC12	2828
IBM zEnterprise EC12	zEC12	2827
IBM zEnterprise 114	z114	2818
IBM zEnterprise 196	z196	2817
IBM System z10® Business Class	z10 BC	2098
IBM System z10 Enterprise Class	z10 EC	2097
IBM System z9® Business Class	z9 BC	2096
IBM System z9 Enterprise Class	z9 EC	2094

The examples in this book use tools, such as the HCD and channel-path identifier (CHPID) Mapping Tool (CMT) that refer to the machine type as opposed to server names. For more information, see Chapter 4, “Mapping CHIDs to CHPIDs by using the CMT” on page 49.

2.3 IBM Resource Link

The first step in planning for the installation of the z14 ZR1 is to access IBM Resource Link. You must register with Resource Link by providing a client site number, ID, and a valid email address. Your IBM representative can assist you with the registration process. After you have an IBM ID, you can customize your profile to accommodate the servers for which you are responsible.

On the Resource Link website, you can access various resources and tools that are designed to help the installation process. Several tools are available to simplify the installation process of a z14 ZR1 server. Even if you worked with most of these tools before, be sure to check for the latest versions that are relevant to z14 ZR1.

The Education and Library tabs on the website display information about the IBM Z family and some online tutorials. Under the Tools tab, you can download the latest version of the most frequently used tools and obtain system and configuration information.

2.4 Hardware Configuration Definition tool

HCD is an application that runs on z/OS and IBM z/VM that supplies an interactive dialog to generate the input/output definition file (IODF) and the input/output configuration data set (IOCDS). Generally, use 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 of the use of HCM (see 2.4.1, “Hardware Configuration Manager” on page 10).

New hardware (z14 ZR1) requires program temporary fixes (PTFs) to enable definition support in HCD.

For the most current information about HCD, see the Hardware Configuration page.

When defining devices in HCD, the hardware features can be selected according to the physical setup of the devices that are attached to the z14 ZR1. Detailed forms and charts that describe the current environment facilitate the planning process.

2.4.1 Hardware Configuration Manager

HCM provides a graphical user interface 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 most updated information about HCM, see the Hardware Configuration page.

2.5 CHPID Mapping Tool

The CMT provides a mechanism to map physical channel IDs (PCHIDs) to CHPIDs as required on a z14 ZR1. The CMT is optional but is preferred to manually mapping the PCHIDs to CHPIDs. The use of the CMT provides the best availability recommendations for a particular configuration.

The following files are needed to obtain an IODF file that contains the correct PCHID numbers by using CMT:

- A production IODF file without PCHID numbers. For more information about how to obtain this file, see Chapter 4, “Mapping CHIDs to CHPIDs by using the CMT” on page 49.
- The CFReport file reflecting the physical configuration of the ordered z14 ZR1 server, which is obtained from the Resource Link website. The CCN is generated by your IBM Client Representative when building the order for your configuration.

2.5.1 HCD and the CMT

The HCD process flow for a new z14 ZR1 installation is shown in Figure 2-1.

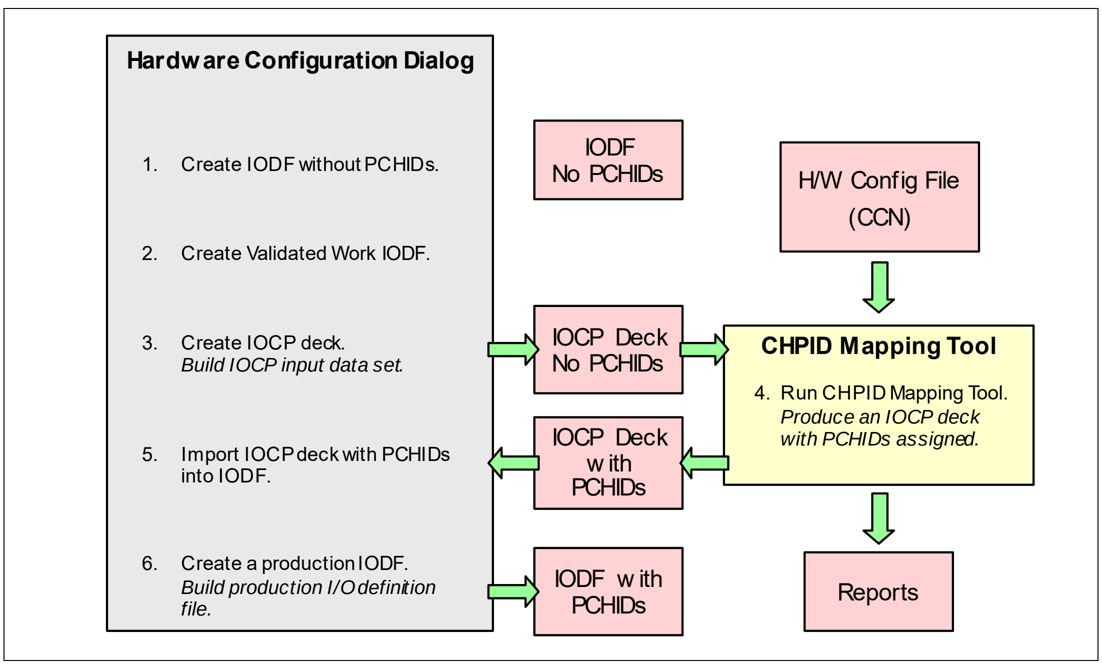

Figure 2-1 CMT: I/O configuration definition flow for a new installation

Part of the tasks that are shown in Figure 2-1 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 site by using a registered Resource Link ID.

For more information, see the *CHPID Mapping Tool Users Guide*, GC28-6984.

For more information about how to use the CMT, see Chapter 4, “Mapping CHIDs to CHPIDs by using the CMT” on page 49.

2.6 Other tools

The tools that are described in this section are not referenced in this book. However, they can help speed up the process of planning and configuring for specific topics that are outside of this book.

2.6.1 Input/output configuration program

ICP IOCP Version 5 Release 4 or later is required for a z14 ZR1 server. You can define the z14 ZR1 configuration by using only IOCP. However, HCD is suggested because of its verification and validation capabilities. By using ICP IOCP, it is possible to write an IOCDS in preparation for a CPC upgrade.

For more information about the changes and requirements for ICP IOCP, see *IBM Z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7172.

2.6.2 World Wide Port Name Prediction Tool

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

One of the most important configuration parameters are WWPNs, which uniquely identify physical or virtual Fibre Channel ports. They are typically used in Fibre Channel 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 is available for download from IBM Resource Link and is applicable to all FICON® channels that are defined as CHPID type FCP (for communication with SCSI devices) on z14 ZR1. For more information about this tool, see this web page (IBMid required).

WWPN Persistence

The FCP WWPNs are determined based on the I/O serial number of the CPC, the IOCDS configuration details (for NPIV WWPNs), and the PCHID values (for physical WWPNs). With the introduction of the z13, the WWPN Persistence configuration option was introduced. When FC 0099 (WWPN Persistence) is ordered as part of a new or upgraded configuration for a z14 ZR1, the I/O serial number part of the WWPN for the new z14 ZR1 is the same serial number as for the source machine configuration.

For more information, see the Techdocs website.

2.6.3 Coupling Facility Structure Sizer

Moving to a new z14 ZR1 means migrating to a higher CFCC level (CFCC level 22). If your existing 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 the amount of storage that must be allocated for coupling facility partitions more accurately. For more information about this tool, see the CFSizer page.

2.6.4 Power Estimation Tool

The Power Estimation Tool is a web-based tool with which you estimate the power consumption for your IBM Z server. The tool also estimates the machine’s weight.

For more information about this tool, see the IBM Resource Link.

2.6.5 Shared Memory Communications Applicability Tool

A tool that is called Shared Memory Communications (SMC) Applicability Tool (SMCAT) was created that helps customers to determine the value of SMC-R and SMC-D in their environment with minimal effort and minimal impact.

SMCAT is integrated within the TCP/IP stack and gathers new statistics that are used to project SMC applicability and benefits for the current system. For more information, see the Shared Memory Communications Reference Information website.

2.6.6 zBNA Tool

zBNA is a PC-based productivity tool that provides a means of estimating the elapsed time for batch jobs solely based on the differences in CPU speeds for a base processor and a target processor, the number of engines on each system, and system capacities. Data sharing is not considered. zBNA provides powerful, graphic demonstration of the z/OS batch window.

The zBNA Tool also provides the capability to project the benefits of deploying the zEDC Express feature and the ability to estimate the benefit of zHyperLink I/O activity.

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

2.7 Hardware Management Console/Support Element setup

This section introduces the configuration and management tools and procedures available on the Hardware Management Console (HMC) and the Support Element (SE).

2.7.1 Defining the HMC Activation Profiles

Activation profiles must be customized by using the HMC. Activation profiles are required for central processor complex (CPC) and CPC image activation. They are used to tailor the operation of a CPC and are stored in the SE that is associated with the CPC. The following types of activation profiles are available:

- 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 CPC previously activated.
- Load: A load profile is used to load an activated image with a control program or operating system.
- Group: A group profile is used to define the group capacity value for all logical partitions belonging to that group.

Default profiles of each of these types are provided. The *Activate* task activates the CPC or CPC image. Initially, the *Default* profile is selected. You can specify an activation profile other than the Default. This feature provides you with the capability to have multiple profiles, for example one for every IOCDS file 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 is loaded and how much main storage is used. Using the reset profile, you must provide the order in which the LPARs are activated during power-on reset (POR). The maximum number of *Reset profiles* that is allowed for each CPC is 26.

Image Profile

Select the appropriate RESET profile and within the profile, select the appropriate IOCDS. The list of LPARs that are defined in the IOCDS is displayed. Parameters must be set for each LPAR before it can be activated and IPLed. The parameters for each LPAR define the following settings:

- General: The mode of operation and its identifier
- Processor: The number of logical CPs, zIIPs, and the weight assigned to the processor
- Security: The security options for this LPAR
- Storage: Memory and Virtual Flash Memory assigned to this LPAR
- Options: The I/O priority and defined capacity options
- Load: The load parameters necessary to IPL this LPAR
- Crypto: The Crypto Express parameters (also see 2.7.2, “Cryptographic configuration” on page 15)

Note: To help you gathering 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 429.

For more information about how to define an Image Profile, see 5.4, “Creating an Image Profile on the 3907 Support Element” on page 99.

Load profile

A Load profile is needed to define the channel address of the device from which the operating system is loaded. Depending on the SE model and machine type, the maximum number of Load profiles that are allowed for each CPC is 511.

Group profile

A Group profile defines the group capacity value that can be customized in determining the allocation and management of processor resources that are assigned to the logical partition in a group.

2.7.2 Cryptographic configuration

The activation profile that you use to activate a logical partition prepares it for running software products that use the Crypto Express feature. The use of the feature’s cryptographic facilities and functions requires customizing the logical partition’s activation profile to complete the following tasks:

- Install the CP Assist for Cryptographic Facility (CPACF) DES/TDES Enablement feature if you are planning to use ICSF.
- Provide it 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 it with an operating system, such as z/OS, that supports the use of cryptographic functions.

2.7.3 Defining the LPAR Group Control

The following methods can be used to limit the processor capacity usage for a group of LPARs and help you control software cost:

- Group Capacity is capping the processor consumption to the value of the four-hour rolling average (4HRA) for a group of LPARs.
- LPAR group absolute capping value is independent of the four-hour rolling average consumption and limits the amount of physical processor capacity that is used by a group of LPARs.

Both of these methods can be used concurrently and in combination with LPAR capping.

Consider reevaluating the parameters in a scenario where the values must be migrated from a previous generation CPC to a z14 ZR1 so that they fit to the new CPC.

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

2.7.4 Defining the Console (HMC part)

The OSA-ICC function of the OSA-Express 1000Base-T feature supports TN3270 enhancements (TN3270E) and non-SNA distributed function terminal (DFT) 3270 emulation. Planning for an IBM z14 Model ZR1 OSA-ICC implementation requires input from the following disciplines within a customer organization:

- IBM Z server I/O subsystem configuration
- Operating system configuration
- OSA-Express feature configuration
- Ethernet LAN configuration
- Client TN3270E configuration

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

- OSA-ICC server: Name, Host IP address, TCP port number, Gateway IP address, the netmask, the network type, and the MTU size
- OSA-ICC session definitions: Channel subsystem, the MIF (LPAR) ID, Device number, LU-name, clients' IP address, clients' DHDTO/RSP/RTO

Note: Consider defining multiple sessions per LPAR to allow access for a number of users at the same time.

For an upgrade of an IBM Z server to a z14 ZR1, these definitions can be exported from the source machine by using on-board HMC facilities and imported back again after the upgrade is complete.

For more information about the definitions, see Chapter 7, "Defining console communication" on page 143. For more information about implementation, see *OSA-Express Integrated Console Controller Implementation Guide*, SG24-6364.

2.7.5 Support Element settings

The SEs that are supplied with the z14 ZR1 are two appliances based on 1U x86 servers. Both units are installed at the top of the A frame. One is a primary SE and the other is the alternative SE.

Generally, the SE settings are considered part of the physical installation of the z14 ZR1 server and not presented in this book.

For a new z14 ZR1 server, a new range of TCP/IP addresses must be provided by the customer to the system services representative (SSR) who performs the physical installation. As an extra measure of security, provisioning a separate LAN segment for the management functions is preferred. During an upgrade from an older IBM Z server to a z14 ZR1, the current settings on the SEs should be backed up for migration purposes.

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

2.7.6 Setting up Server Time Protocol

STP provides the means by which the time of day (TOD) clocks in various systems can be synchronized by using messages that are transported over coupling links. STP operates along with the TOD-clock steering facility, which provides a new timing mode, timing states, external interrupts, and machine check conditions.

> **STP connectivity for z14 ZR1 and CTN roles:** The z14 ZR1 server does not support coupling connectivity by using the InfiniBand feature. As such, the z14 ZR1 CPC only can connect for transmitting coupling or timing (STP) data to a z13/z13s or to another z14 M0x/z14 ZR1 CPC. In a CTN that also contains zEC12/zBC12 servers, z14 ZR1 cannot play a role in the CTN (PTS/BTS/Arbiter) for availability reasons.

The HMC provides the user interface to manage an STP-only Coordinated Timing Network (CTN).

Consider the following points when setting up an HMC for STP:

- A CTN ID must be unique for all IBM Z servers that will be part of the CTN.
- To synchronize IBM Z servers to an External Time Source (ETS), network Time Protocol (NTP) server information (and network connectivity that uses NTP/NTPS protocol with optional pulse per second [PPS]) must be provided.
- Customer must have the time zone offset, Daylight Saving Time offset, and leap second offset.
- Optional, the HMC can be configured as an NTP server.
- For the IBM Z servers that are part of a CTN, STP roles must be planned (Preferred, Backup, and Current Time Servers and Arbiter).
- As part of a migration, changing the Current Time Server must be done before migration to the new platform (z14 ZR1).

> **Note:** The z14 ZR1 supports STP stratum level 4. This feature avoids the added complexity and expense of system reconfiguration. This change must be installed all systems that might become exposed to this situation. Stratum level 4 should be used only during a migration, and for a short period.

For more information, see Chapter 8, “Preparing for Sysplex and configuring Server Time Protocol” on page 159.

For more information about planning, implementing, and managing an STP environment, see the following publications:

- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Recovery Guide*, SG24-7380

2.8 Activities centered on the IODF

This section describes the information (I/O configuration) in the IODF.

2.8.1 Logical channel subsystems

An IBM Z processor manages I/O resources (including logical partitions, channel paths, control units, and I/O devices) by housing them in multiple logical channel subsystems. Each logical channel subsystem (LCSS) can have up to 256 channel paths. The z14 ZR1 supports up to 3 LCSSs.

A spanned channel path is one that can be used by partitions in more than one logical channel subsystem. You must use the same CHPID value across all logical channel subsystems that share a spanned channel. However, logical channel subsystems 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.

Consider the use of multiple logical channel subsystems during the planning phase. By using multiple logical channel subsystems, you can logically partition your physical channel resources to accommodate large-scale enterprise workload connectivity and high-bandwidth demands.

Each LCSS can have up to 256 CHPIDs. On the z14 ZR1, you can define up to three LCSSs. Each LCSS can support up to 15 logical partitions (LPARs) except for LCSS 2, which can support up to 10 LPARs for a total of 40 LPARs per z14 ZR1 server.

Also, LCSSs provide for multiple subchannel sets for expanding the number of I/O devices that are managed in each CSS. The z14 ZR1 supports up to three subchannel sets per LCSS.

Not all device types are eligible for nonzero subchannel sets. Subchannel set 0 (SS0) can be used for any type of device. More subchannel sets (for example: subchannel set 1 [SS1]) can be used for certain classes of devices only, such as parallel access volume alias devices.

For more information, see *IBM z14 Model ZR1 Technical Guide*, SG24-8651. Use multiple subchannel sets to move devices of eligible device types to extra subchannel sets, then define more physical devices to SS0.

2.8.2 Defining partitions

The IBM Processor Resource/System Manager (PR/SM) feature allows a single CPC to run multiple operating systems in LPAR mode. Each operating system has its own logical partition, which is a separate set of system resources that includes the following items:

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

Only LPAR mode (not basic mode) is supported on IBM Z servers.

Profile data can be exported on the older server and imported on the z14 ZR1. If the LPAR data was imported from an older server, consider the LPAR sizing before the LPAR migration to the z14 ZR1. For more information, see the IBM Resource Link (log in required).

For more information about how to define LPARs in IODF, see Chapter 3, “Preparing for a new z14 ZR1” on page 31.

Planning considerations for Virtual Flash Memory

IBM Virtual Flash Memory (VFM - Feature Code 0614) is the replacement for the Flash Express features (FC 0402 and FC 0403).

IBM VFM includes the following minimum software requirements:

- z/OS V2.3.
- z/OS V2.2.
- z/OS V2.1.
- z/OS V1.13 with PTFs, the z/OS V1.13 RSM Enablement Offering web deliverable installed, and an extended support contract for IBM Software Support Services. The web deliverable is available at the z/OS downloads page.

VFM (FC 0614) is available in 512 GB increments, each feature providing for 512 GB of memory. Up to four VFM features can be ordered, which results in a total of 2 TB of virtual flash memory. The plan ahead memory option must consider VFM requirements.

With the introduction of VFM, the existing operating system interface is *not* changed to handling the storage-class memory (SCM). Operating systems handle VFM the same way as the Flash Express. The allocation of VFM storage is done during LPAR activation because the LPAR hypervisor manages the partition memory.

The initial and maximum amounts of VFM are specified in the LPAR image profile. VFM can be added or deleted to or from operating systems by using 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.

Virtual Flash Memory 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 such, it is recommended to assign the maximum amount installable (2 TB) for all LPARs that are candidates for the use of VFM and set initial allocation to 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 allows more storage to be added to partitions subject to the amount that is not assigned to other partitions. For more information, see 10.3.3, “Configuring VFM” on page 247.

If the total amount of VFM that is allocated to all active partitions is equal to the LICCC value, but the sum of active partition maximums is larger than the installed amount, a customer might concurrently add VFM and increase allocations without reactivating partitions. This feature is shown in the examples that are described next.

Non-disruptive migration

An example of a non-disruptive migration includes the following features:

- A z14 ZR1 CPC has three VFM features installed (512 GB each), LICCC = 1.5 TB.
- LPAR A has 1.0 TB assigned, max = 1.5 TB.
- LPAR B has 512 GB assigned, max = 1.0 TB.
- LPAR B must be altered to have 1.0 TB assigned. This change is not possible within the constraints of the installed VFM.

- Another 512 GB VFM feature is purchased and installed concurrently. Now up to 512 GB can be added concurrently to LPAR B without reactivating the LPAR.

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

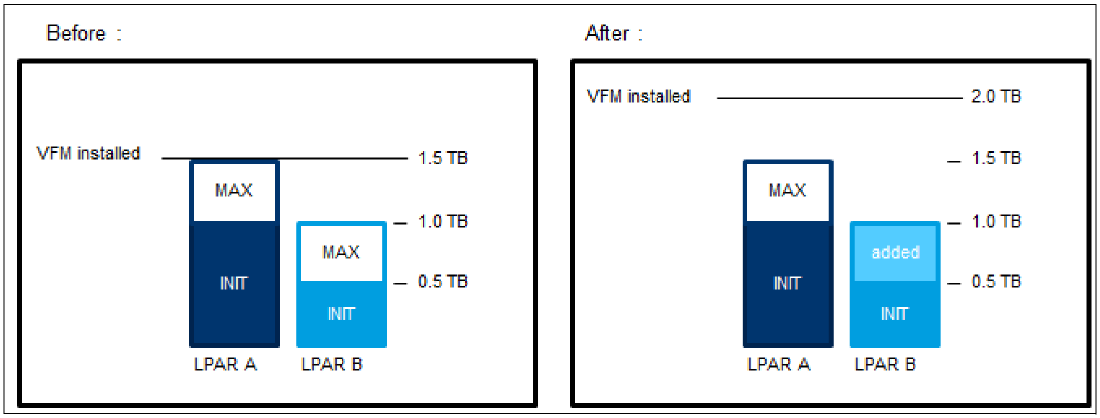

Figure 2-2 Non-disruptive VFM migration example

Disruptive migration

An example of a disruptive migration includes the following features:

- A z14 ZR1 CPC has two VFM features installed (512 GB per feature), LICCC = 1.0 TB.
- LPAR A has 512 GB TB assigned, max = 1.0 TB.
- LPAR B has 256 GB assigned, max = 1.0 TB.
- LPAR A must be altered to have up to 1.5 TB. This change falls outside the range of maximum installed VFM.
- Two extra 512 GB VFM features are purchased and activated concurrently (assuming plan ahead memory was ordered and memory is available). LPAR A must be reactivated with the new maximum VFM value of at least 1.5 TB and less than or equal to 2.0 TB.

Figure 2-3 shows the disruptive migration example.

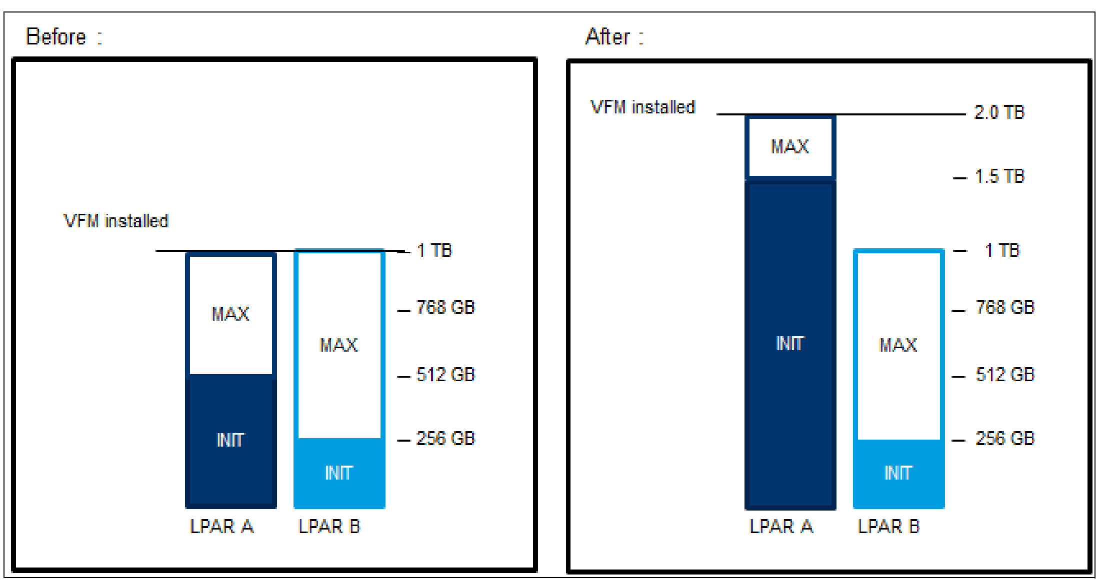

Figure 2-3 Disruptive VFM migration example

For more information about how to configure VFM, see 10.3, "Virtual Flash Memory" on page 247.

2.8.3 Defining Storage I/O - FICON and FCP

FICON Express16S+, FICON Express 16S, and FICON Express8S features provide connectivity to storage devices by using Fibre Connection (FICON) or Fibre Channel Protocol (FCP). FICON Express16S+ and FICON Express 16S features support auto negotiation for the link data rate: 4 Gbps, 8 Gbps, and 16 Gbps. FICON Express8S supports auto negotiation for the link data rate at 2 Gbps, 4 Gbps, and 8 Gbps.

FICON Express16S+, FICON Express16S, and FICON Express8S support High-Performance FICON for z IBM z Systems (zHPF). zHPF is an extension to the FICON architecture that provides performance improvement for single-track and multi-track operations.

On a new build z14 ZR1 server, only the FICON Express16S+ feature can be ordered. The FICON Express16S and FICON Express8S features can be carried forward when upgrading from an older IBM Z server.

> **Note:** On a FICON Express16S+ feature, *both* ports must be configured as channel type FC or FCP. A mixed configuration is *not* allowed.

For more information about how to configure a FICON Express16S+ feature, see Chapter 12, "Adding storage devices" on page 275.

2.8.4 Defining the IBM zHyperLink Express

For more information about defining zHyperLink Express, see 10.6, "IBM zHyperlink Express" on page 262.

> **Important:** IBM intends to deliver IMS exploitation of IBM z14 and DS8880 zHyperLink WRITE operations[a]. zHyperLink Express is a direct connect short distance IBM Z I/O adapter that is designed to work with a FICON or High-Performance FICON SAN infrastructure.

a. IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

2.8.5 Defining Network

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

- Open Systems Adapter (OSA)
- Shared Memory Communications (SMC-R and SMC-D):
  - SMC - RDMA over Converged Ethernet (RoCE) Express features (SMC-R)
  - SMC - Direct Memory Access over Internal Shared Memory (SMC-D)
- HiperSockets

Open Systems Adapter

The OSA Express features are installed in an IBM z14 Model ZR1server PCIe+ I/O drawer. The features are available as different types and support several networking protocols. Depending on the types of OSAs installed in the z14 ZR1, the CPC supports attachment with the following characteristics:

- Copper-based Ethernet (10, 100 and 1000 Mbps)
- Fiber-based Gigabit Ethernet (GbE), Short Wave (SX), and Long Wave (LX)
- Fiber-based 10-Gigabit Ethernet Short Reach (SR) and Long Reach (LR)

Based on the intended use, the operating modes must be defined with channel type and device address. For more configuration information, see Chapter 6, “Configuring network features” on page 123 and the *OSA-Express Implementation Guide*, SG24-5948.

Starting with Driver Level 22 (HMC 2.13.0) installed on z13, HMC was enhanced to take advantage of the Open Systems Adapter/Support Facility (OSA/SF) function for the OSA-Express6S, OSA-Express5S, and OSA-Express4S features. OSA/SF on the HMC or the OSA/SF in the operating system component can be used for the OSA-Express4S features. For the OSA-Express6S and OSA-Express5S features, OSA/SF on the HMC is required. The OSA/SF is used primarily for the following purposes:

- Manage all OSA ports.
- Configure all OSA non-QDIO ports.
- Configure local MAC addresses.
- Display registered IPv4 addresses (in use and not in use). It is supported on IBM Z platform for QDIO ports.
- Display registered IPv4 or IPv6 Virtual MAC and VLAN ID associated with all OSA Ethernet features configured as QDIO Layer 2.
- Provide status information about an OSA port and its shared or exclusive use state.

For more information about the use of OSA/SF on the HMC, see 6.3, “Customizing OSA-Express using OSA Advanced facilities” on page 127.

OSA-Express6S 1000BASE-T adapters[a]: OSA-Express6S 1000BASE-T adapters (FC 0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mbps link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mbps (1 Gbps) link speed.

a. IBM’s statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM’s sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

Shared Memory Communications - RDMA

The 10GbE RoCE Express (FC 0411) and 10GbE RoCE Express2 (FC 0412) features are designed to help reduce CPU consumption for applications that use the TCP/IP stack without requiring application changes. The use of the RoCE Express features also helps to reduce network latency by using the SMC-R protocol in z/OS V2.1 or later. For more information, see RFC 7609. SMC-R is transparent to applications and can be used for LPAR-to-LPAR communications on a single CPC or for server-to-server communications across multiple IBM Z CPCs.

Deployment of the RoCE Express features is supported in a point-to-point configuration or switched configurations. When planning to deploy RoCE Express 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) disabled
- No firewalls, no routing

IBM provides the SMC Applicability Tool (SMCAT) that helps determine the potential gains of using SMC-R in an environment (see 2.6.5, “Shared Memory Communications Applicability Tool” on page 13).

With z14 ZR1, the new 10GbE RoCE Express2 feature is available. This feature provides increased virtualization (sharing) capabilities. For more information, see *IBM z14 Model ZR1Technical Guide*, SG24-8651.

RoCE Express features port configuration: Consider the following points:

- For 10GbE RoCE Express2 feature (FC 0412), the port number is now configured with the FID number in HCD (or IOCDS) and Port number *must* be configured (*no default exists*).
- Port number for 10GbE RoCE Express (FC 0411) is configured in z/OS TCP/IP profile and does not change.

When defining a FID in the TCP/IP profile for 10GbE RoCE Express2 (FC 0412), the port number is no longer applicable.

When preparing to deploy the RoCE Express features, consider the following items:

- The RoCE Express features are “Native” PCIe features; therefore, the following configuration items must be provided:
  - Function ID
  - Type
  - PCHID
  - Virtual Function ID (VF)
  - Port number
- Determine which LPARs are to be shared by one 10GbE RoCE Express port.
- Assign the VFs between the sharing LPARs.

For more configuration information, see 15.2.3, “Defining a RoCE-2 PCIe function” on page 375.

For 10GbE RoCE Express2 feature management information, see 10.4.4, “SMC-R Management” on page 259.

Consider Native PCIe feature Plugging and Resource Groups

The native PCIe feature support is provided by Resource Group (RG) code running on the integrated firmware processor (IFP). For resilience, four independent RGs are always on the system that share the IFP. For high availability purposes, always use at least two PCIe features located in different RGs, as shown in Figure 2-4.

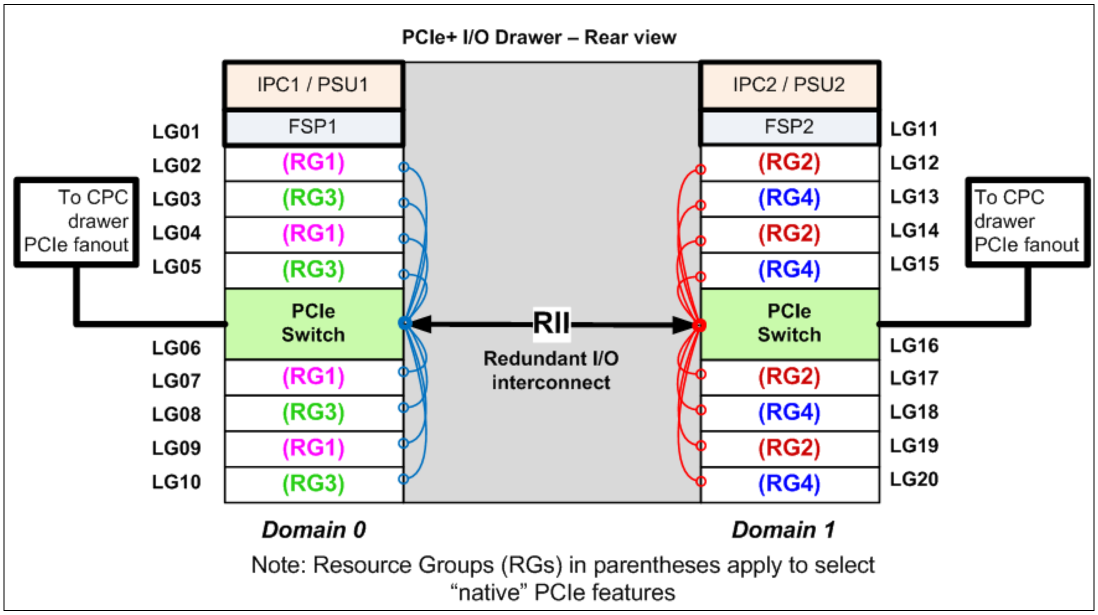

Figure 2-4 Relationship among PCIe+ I/O drawer slots, domains, and RGs in the z14 ZR1

Shared Memory Communications - Direct Memory Access (SMC-D)

With the z13 (Driver 27) and z13s servers, IBM introduced SMC-D. SMC-D uses Internal Shared Memory (ISM) virtual PCIe adapter to provide direct memory access communications between LPARs inside the same IBM Z CPC.

SMC-D maintains the socket-API transparency aspect of SMC-R so that applications that use TCP/IP communications can benefit immediately without requiring any application software or IP topology changes. SMC-D completes the overall Shared Memory Communications solution, which provides synergy with SMC-R. Both protocols use shared memory architectural concepts, which eliminates TCP/IP processing in the data path, yet preserves TCP/IP Qualities of Service for connection management purposes.

From a planning standpoint, SMC-D is similar to SMC-R; therefore, the same planning considerations apply. The objective is to provide consistent operations and management tasks for SMC-D and SMC-R. SMC-D uses a new virtual PCI adapter that is called *Internal Shared Memory* (ISM). The ISM Interfaces are associated with IP interfaces; for example, HiperSockets or OSA. ISM interfaces do not exist without an IP interface.

ISM interfaces are not defined in software. Instead, ISM interfaces are dynamically defined and created, and automatically started and stopped. You do not need to operate (Start or Stop) the ISM interfaces. Unlike RoCE, ISM FIDs (PFIDs) are not defined in software. Instead, they are auto-discovered based on their PNet ID.

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

- Run the SMCAT to evaluate applicability and potential value. For more information about the SMCAT, see the IBM z/OS SMC Applicability Test (SMC-AT) document.

  For more information, see 2.6.5, "Shared Memory Communications Applicability Tool" on page 13.

- Review and adjust as needed the available real memory and fixed memory usage limits (z/OS and CS). SMC requires fixed memory. You might need to review the limits and provision extra real memory for z/OS.
- Review IP topology, VLAN usage considerations, and IPSec.

- Review changes to messages, monitoring information, and diagnostic tools. Many updates are available for the following items:
  - Messages (IBM VTAM® and TCP stack)
  - **`Netstat`** (status, monitoring, and display information)
  - CS diagnostic tools (VIT, Packet trace, CTRACE, and IPCS formatted memory dumps)

For more information about SMC-R planning and security considerations, see the SMC-R tab on the Shared Memory Communications Reference Information page.

For more information about SMC-D planning and security considerations, see the SMC-D tab on the Shared Memory Communications Reference Information page.

For more information about how to define SMC-D, see 15.2.2, "Defining an ISM PCIe function" on page 372.

For an overview of how to manage an SMC-D connection, see 10.5.4, "SMC-D management" on page 261.

HiperSockets

HiperSockets provides the fastest TCP/IP communications between z/OS, z/VM, IBM z/VSE®, and Linux logical partitions within a z14 ZR1 CPC, that act like internal "virtual" local area networks. This HiperSockets implementation is achieved by using the Licensed Internal Code (LIC) and supporting device drivers in the operating systems. HiperSockets establish a network with higher availability, security, simplicity, performance, and cost effectiveness than can be achieved by using an external IP network.

The HiperSockets function is based on the OSA-Express queued direct input/output (QDIO) protocol and therefore, HiperSockets is called internal QDIO (iQDIO). The 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, which bypasses external network delays and provides users high-speed logical LANs with minimal system and network overhead.

HiperSockets can be defined as Multiple Image Facility (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 logical partitions you want to grant access to that LAN.

HiperSockets is supported by the following operating systems:

- All in-service z/OS releases
- All in-service z/VM releases
- All in service z/VSE releases
- Linux on Z

On a z14 ZR1, HiperSockets supports the following functions:

- HiperSockets Broadcast

  Supported across HiperSockets on Internet Protocol Version 4 (IPv4) for applications. Applications that use the broadcast function can propagate the broadcast frames to all TCP/IP applications that use HiperSockets. This support is applicable in Linux, z/OS, and z/VM environments.

- VLAN support

  Virtual local area networks (VLANs) are supported by Linux on z Systems and z/OS V1R8 or later for HiperSockets. VLANs can reduce overhead by allowing networks to be organized by traffic patterns rather than physical location. This enhancement allows traffic flow on a VLAN connection over HiperSockets and between HiperSockets and OSA-Express Ethernet features.

- IPv6 support on HiperSockets
- HiperSockets Network Concentrator

  Traffic between HiperSockets and OSA-Express can be transparently bridged by using the HiperSockets Network Concentrator. This configuration eliminates intervening network routing overhead, which results in increasing performance and a simplified network configuration. This improvement is achieved by configuring a connector Linux system that has HiperSockets and OSA-Express connections defined to it.

- HiperSockets Layer 2 support

  HiperSockets supports two transport modes on the z14 ZR1 Layer 2 (Link Layer) and Layer 3 (Network and IP Layer).

  As with Layer 3 functions, HiperSockets Layer 2 devices can be configured as primary or secondary connectors or multicast routers. These configurations enable high-performance and highly available Link Layer switches between the HiperSockets network and an external Ethernet.

- HiperSockets multiple write facility

  HiperSockets performance was increased by allowing streaming of bulk data over a HiperSockets link between logical partitions. Multiple writes with fewer I/O interrupts reduce processor usage both the sending and receiving logical partitions, and is supported in z/OS.

- HiperSockets Completion Queue

  The HiperSockets Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible, and asynchronously if necessary. This function combines ultra-low latency with more tolerance for traffic peaks.

  With the asynchronous support, during high volume situations, data can be temporarily held until the receiver has buffers available in its inbound queue. This function provides end-to-end performance improvement for LPAR to LPAR communication.

- HiperSockets Virtual Switch Bridge Support

  The z/VM virtual switch is enhanced to transparently bridge a guest virtual machine network connection on a HiperSockets LAN segment. z/VM 6.2 or later, TCP/IP, and Performance Toolkit APARs are required for this support.

  This bridge allows a single HiperSockets guest virtual machine network connection to also directly communicate with the following devices:

  – Other guest virtual machines on the virtual switch
  – External network hosts through the virtual switch OSA UPLINK port

- zIIP-Assisted HiperSockets for large messages

  In z/OS, HiperSockets was enhanced for zIIP exploitation. Specifically, the z/OS Communications Server allows the HiperSockets Multiple Write Facility processing for large outbound messages that originate from z/OS to be run on a zIIP.

  z/OS application workloads that are based on XML, HTTP, SOAP, Java, and traditional file transfer can benefit from zIIP enablement by lowering general-purpose processor usage.

  When the workload is eligible, the HiperSockets device driver layer processing (write command) is redirected to a zIIP, which unblocks the sending application.

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

2.8.6 Defining the console (OSA-ICC)

The OSA-ICC function of the OSA-Express 1000Base-T feature supports TN3270 enhancements (TN3270E) and non-SNA DFT 3270 emulation. Planning for an IBM Z z14 Model ZR1 OSA-ICC implementation requires input from several disciplines within a customer organization.

The following aspects of system configuration provide input for configuring OSA-ICC:

- IBM Z server I/O subsystem configuration
- Operating system configuration
- OSA-Express feature configuration
- Ethernet LAN configuration
- Client TN3270E configuration

In HCD, the OSA-Express feature must be defined to operate as an Integrated Console Controller (ICC). The configuration includes the following requirements:

- IBM Z server I/O subsystem configuration: The same basic rules for adding an OSA-ICC adapter apply as to any other new device.
- Operating system configuration: To have a Nucleus Initialization Program (NIP) console available, ensure that the correct device number is defined in the HCD Operating system *Work with consoles* dialog.

During an upgrade from an IBM Z server to a z14 ZR1, the same definitions can be used for the new machine as on the source configuration.

For more implementation information, see *OSA-Express Integrated Console Controller Implementation Guide*, SG24-6364.

The following planning topics must be considered:

- Reserve at least one OSA-Express 1000Base-T port to be defined as channel type OSC
- Define 3270-X Devices in HCD to act as system consoles
- The use of OSA/Advanced facilities to configure the sessions

For more information about how to configure non-SNA consoles, see Chapter 7, “Defining console communication” on page 143.

2.8.7 Defining coupling and timing only links

Support for Parallel Sysplex includes the Coupling Facility Control Code and coupling links. A new Coupling connectivity in support of Parallel Sysplex environments is provided on the z14 ZR1 by the following features:

- Coupling Express Long Reach (CE LR). The feature (FC 0433) has two ports coupling link connectivity for a distance up to 10 km (6.2 miles).
- Integrated Coupling Adapter (ICA SR), which is FC 0172.
- Internal Coupling (ICs) channels operate at memory speed.

For more information, see *IBM Z Connectivity Handbook*, SG24-5444.

All coupling link types can be used to carry STP messages.

Note: The CE LR is a two-port card that occupies one PCIe+ I/O drawer slot. Therefore, an IBM z14 Model ZR1server that is configured as a stand-alone Coupling Facility (CF) must include at least one PCIe+ I/O drawer.

Planning considerations

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

Coupling connectivity for z14 ZR1: The z14 ZR1 CPC *does not support coupling connectivity using InfiniBand features*. As such, it can connect only for transmitting coupling or timing (STP) data to a z13/z13s or to another z14 M0x/z14 ZR1 CPC.

In a Parallel Sysplex that also contains zEC12/zBC12 servers, the z14 ZR1 or the zEC12/zBC12 *cannot* be used for running the Coupling Facility LPAR. The CF LPAR must be run on a CPC that includes coupling connectivity to z14 ZR1 and the zEC12/zBC12.

Depending on the hardware that is configured on the CPC, a different channel type must be defined.

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 Coupling Express Long Reach links, dense wavelength division multiplexing (DWDM) technology can be used to extend the maximum length of the CF links.

For more information about qualified devices, see the IBM Resource Link.

STP signals can be exchanged between two CPCs without any CF LPARs involved. If physical coupling links are established between two CPCs, HCD allows 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 Sysplex and configuring Server Time Protocol" on page 159.

IBM z14 Model M0x (machine type 3906) will be the last z Systems and IBM Z server to support HCA3-O and HCA3-O LR adapters[a]: z14 M0x will be last z Systems and IBM Z server to support HCA3-O fanout for 12x IFB (#0171) and HCA3-O LR fanout for 1x IFB (#0170). As announced previously, z13s is the last mid-range z Systems server to support these adapters. Enterprises should begin migrating from HCA3-O and HCA3-O LR adapters to ICA SR and Coupling Express Long Reach (CE LR) adapters on z14, z13, and z13s.

For high-speed short-range coupling connectivity, enterprises should migrate to the Integrated Coupling Adapter (ICA-SR). For long-range coupling connectivity, enterprises should migrate to the new Coupling Express LR coupling adapter. For long-range coupling connectivity requiring a DWDM, enterprises must determine their wanted DWDM vendor's plan to qualify the planned replacement long-range coupling link.

a. IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

2.8.8 Planning considerations for zEDC

This section provides planning considerations for installing the zEDC Express feature in a z14 ZR1.

The zEDC Express feature is a hardware feature that allows for data compression and decompression. It is a PCIe native feature that allows for high-performance, low-latency compression that reduces processor use. The hardware device is a standard computer expansion card that is installed into the PCIe I/O drawer.

Be sure to install a minimum of two zEDC Express features, one per Resource Group (RG). For the best data throughput and availability (two features per RG) for a total of four features, must be installed. For the full zEDC benefit, zEDC should be active on all systems that might access or share compressed format data sets. This configuration eliminates instances where software inflation is used when zEDC is not available.

A more information about the zEDC Express feature, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

In this section, a short summary about planning consideration of the zEDC Express feature is given. Several tasks must be completed to use zEDC features:

1. Planning the installation:
   - Consider the number and sharing of one or more zEDC Express features.
   - Update the IFAPRDxx PARMLIB member in z/OS 2.1.
   - Plan for IPLs before activating the prized 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.
4. HCD: Defining the PCIe features:

   For more information, see Chapter 10, “Specialized features” on page 215.
5. Managing the zEDC Express PCIe features:

   For more information, see 10.2.4, “Handling zEDC” on page 244.

3

Preparing for a new z14 ZR1

This chapter describes the following scenarios when preparing for a z14 ZR1 server installation:

- Upgrading an existing IBM Z server (z13s) to a z14 ZR1 server and maintaining an existing serial number.
- Installing a new z14 ZR1 server.

Our examples show upgrading from an existing z13s or installing a new z14 ZR1 only. An upgrade includes a new frame, new drawers (CPC and PCIe+ I/O drawers), and new or carry-forward I/O features.

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

This chapter includes the following topics:

3.1 Hardware features

This section describes the channel (CHPID) types and hardware features for a z14 ZR1.

CHPID type HYL is new for 3907.

The following hardware features are new or available for order with a 3907:

- FC #0427 - FICON Express16S+ LX
- FC #0428 - FICON Express16S+ SX
- FC #0422 - OSA-Express6S GbE LX
- FC #0423 - OSA-Express6S GbE SX
- FC #0424 - OSA-Express6S 10 GbE LR
- FC #0425 - OSA-Express6S 10 GbE SR
- FC #0426 - OSA-Express6S 1000BASE-T Ethernet
- FC #0412 - RoCE Express2 10 GbE
- FC #0893 - Crypto Express6S
- FC #0433 - Coupling Express LR
- FC #0431 - zHyperLink Express
- FC #0614 - IBM Virtual Flash Memory (VFM)
- FC #0172 - Integrated Coupling Adapter SR
- FC #0901 - Regional Crypto Enablement (RCE)

The following CHPID types are migrated (carry forward):

- FC and FCP
- OSC, OSD, OSM, and OSX
- CL5, CS5, and ICP
- IQD

The following hardware features can be migrated to the 3907:

- FC #0418 - FICON Express16S LX
- FC #0419 - FICON Express16S SX
- FC #0409 - FICON Express8S LX
- FC #0410 - FICON Express8S SX
- FC #0413 - OSA-Express5S GbE LX
- FC #0414 - OSA-Express5S GbE SX
- FC #0415 - OSA-Express5S 10 GbE LR
- FC #0416 - OSA-Express5S 10 GbE SR
- FC #0417 - OSA-Express5S 1000BASE-T Ethernet
- FC #0411 - RoCE Express 10 Gbe SR
- FC #0420 - zEDC Express
- FC #0890 - Crypto Express5S
- FC #0172 - Integrated Coupling Adapter SR
- FC #0901 - Regional Crypto Enablement (RCE)

The following CHPID types are not migrated (no carry forward):

- CIB
- OSN

The following hardware features cannot be ordered nor carried forward for an upgrade to 3907:

- FC #3325 - FICON Express8 LX
- FC #3326 - FICON Express8 SX
- FC #0865 - Crypto Express4S

- FC #0403 - Flash Express
- FC #0402 - Flash Express
- FC #0170 - HCA3-O 1x LR IFB
- FC #0171 - HCA3-O 12x IFB
- IBM zAware Firmware
- STP Mixed CTN

3.2 Upgrading a z13s to a z14 ZR1: Maintaining existing serial number (Frame roll MES)

This section describes the steps to upgrade a z13s server that is defined in your IODF to a z14 ZR1 server and maintaining the system serial number.

3.2.1 Scenario overview

This scenario describes the configuration steps to upgrade a 2965 (z13s) CPC to a 3907 (z14 ZR1) CPC. Consider the following key factors:

- Hardware configuration definition (HCD) requires a new CPC (processor) ID for the 3907.
- Keep the same CPC name for the 3907 (this is optional, the CPC name can be changed).
- The 3907 processor channels connect to the same switch ports and access the same control unit interfaces.
- The control unit interfaces connect to the same switch ports.
- The starting IODF is the current 2965 *production* I/O definition file (IODF).
- The target IODF is a new 3907 *work* IODF.
- HCD actions:
  - Migrate updated IOCP statements.
  - Build production IODF.
  - Remote write IODF to input/output configuration data set (IOCDS).
- The HMC actions:
  - Build the Reset Profile and point to required IOCDS.
  - Build/verify the Image Profiles.
  - Build/verify the Load Profiles.
  - Perform a power-on reset.

The example uses a 2965-N20 with a Processor ID of *LEPUS* with three CSSs (CSS ID=0 to CSS ID=2). This system is replaced with a 3907-ZR1 with a Processor ID of *MUSCA2* and three CSSes.

The CPC name LEPUS and serial number are not changed.

The migration options and tool requirements are listed in Table 3-1. For more information about the process steps, see “HCD: Migrating the existing 2965 IODF” on page 34.

Table 3-1 2965 I/O configuration migrated to a 3907

2965 to 3907	Upgrade existing 2965 to a 3907 (MES upgrade)
Processor ID	Required to change the Processor ID to a new ID
CPC name	Generally should be the same name

2965 to 3907	Upgrade existing 2965 to a 3907 (MES upgrade)
Channel to switch port connections	Same ports
Control Unit to switch port connections	Same ports
Starting IODF	Current active *production* IODF
Target IODF	Create a *work* IODF
HCD action	Repeat and change
CHPID Mapping Tool Program	Optional, but good for verifying configuration
CFReport file (CCN)	Required for CMT
IOCP (import from validated work IODF)	Yes
CHPID Mapping Tool actions (PCHID reset)	Yes
CHPID Mapping Tool IOCP Output	Yes
CHPID Mapping Tool Reports	Yes, CHPID and CHPID to CU Report

HCD: Migrating the existing 2965 IODF

The following steps describe how to upgrade an existing 2965 Processor in your IODF to the new 3907 Processor using HCD. Then, migrate the I/O configuration and logical partitions from the 2965 to the 3907. Using HCD, the sequence includes the following steps:

1. Creating the work IODF from the current 2965 production IODF.
2. Repeating the 2965 processor to be replaced.
3. Coupling Link information messages.
4. Deleting any unsupported items in the repeated 2965.
5. Changing the M/T 2965 to 3907 and deleting the 2965.
6. Deleting the 2965 processor definition.
7. Reconnecting the CF channel paths that were not migrated.
8. Using OSA/SF to save and restore OSE OAT configuration data.
9. Exporting and importing OSA-ICC configuration data with OSA Advanced Facilities.
10. Using OSA Advanced Facilities to set OSA parameters.

3.2.2 Creating the work IODF from the current 2965 production IODF

HCD is the tool that is used to make a work IODF, but first we start from the current production IODF that contains the 2965 processor we are upgrading (in our example, `SYS6.IODF78`).

3.2.3 Repeating the 2965 processor to be replaced

To repeat the 2965 processor in HCD, complete the following steps:

1. From the main HCD panel, select option **1.3. Processor List**.
2. In the Processor List (see Figure 3-1 on page 35), enter `r` (for repeat) next to the 2965 that you want to upgrade 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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
r LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
```

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, `SYS6.IODF79.WORK`).
5. The Repeat Processor panel opens (see Figure 3-2). Enter the Processor ID of the new 3907 (in this example, `MUSCB`), keep all the other fields unchanged, and press Enter.

```
*----------------------------- Repeat Processor ------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . MUSCB___                                     |
|                                                                             |
| Processor type . . . . . . . : 2965                                         |
| Processor model  . . . . . . : N20                                          |
| Configuration mode . . . . . : LPAR                                         |
|                                                                             |
| Serial number  . . . . . . . . 0BB4B72965                                   |
| Description  . . . . . . . . . Lepus                                        |
|                                                                             |
| Specify SNA address only if part of an System z cluster:                    |
|                                                                             |
| Network name . . . . . . . . . IBM390PS  +                                  |
| CPC name . . . . . . . . . . . LEPUS     +                                  |
|                                                                             |
| Local system name  . . . . . . LEPUS                                        |
|                                                                             |
*------------------------------------*                                        |
| New IODF SYS6.IODF79.WORK defined. |                                        |
*------------------------------------* ---------------------------------------*
```

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

3.2.4 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 in the example because the CF Link CHPIDs were not copied to the 3907 definition.

```
*------------------------------- Message List --------------------------------*
|   Save  Query  Help                                                         |
| --------------------------------------------------------------------------- |
|                                                               Row 24 of 37 |
| Command ===> ___________________________________________ Scroll ===> CSR   |
|                                                                             |
| Messages are sorted by severity. Select one or more, then press Enter.     |
|                                                                             |
| / Sev Msg. ID  Message Text                                                 |
| #              processor MUSCA is not copied.                               |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              1.F4 of processor LEPUS and channel path 0.F5 of             |
| #              processor LEPUS is not copied.                               |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              1.F5 of processor LEPUS and channel path 0.F4 of             |
| #              processor LEPUS is not copied.                               |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              1.F6 of processor LEPUS and channel path 0.F7 of             |
| #              processor LEPUS is not copied.                               |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              1.F7 of processor LEPUS and channel path 0.F6 of             |
| #              processor LEPUS is not copied.                               |
| _ I   CBDG271I Requested action on object LEPUS successfully processed.    |
*-----------------------------------------------------------------------------*
```

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 LEPUS successfully processed` message.
2. Press PF3 or PF12 to continue. As shown in Figure 3-4, an extra 2965 processor named MUSCB is available.

```
                              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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
_ MUSCB    2965     N20      LPAR  0BB4B72965 Lepus
```

Figure 3-4 Processor List: Repeated processor

3.2.5 Deleting any unsupported items in the repeated 2965

If you are upgrading a processor that contains any CHPID types of OSN or CIB, they must be deleted from the IODF before changing the processor type to 3907.

Note: z14 ZR1 machine type 3907 does *not* support any InfiniBand coupling links. Therefore, all CHPIDs of type CIB must be deleted in an IODF for machine type 3907 and, if necessary, replaced with CHPID types CS5 or CL5.

To delete no longer supported CHPIDs, complete the following steps:

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

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
s MUSCB    2965     N20      LPAR  0BB4B72965 Lepus
```

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 3 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . : MUSCB        Lepus

  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     14692   65535     480     65535     0       0         0
_ 1   65280     95      65535     0       65535     0       0         0
_ 2   65280     0       65535     0       65535     0       0         0
```

Figure 3-6 Channel Subsystem List: Selected CSS

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

```
  Goto  Filter  Backup  Query  Help
----- *-------------------------------------* ------------------------------
      | 1  1. Set Filter                    | t      Row 1 of 46 More:
Comma |    2. Clear Filter                  | ______________ Scroll ===> CSR
      |    3. Count rows on (filtered) list |
Selec *-------------------------------------* nter. To add use F11.

Processor ID . . . . : MUSCB       Lepus
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 11    7E0   IQD   SHR   __  __ __       No  IQD for ISM
_ 28    1E4   FC    SHR   02  02 13       No  fc sw - san64b-b-p19
_ 29    120   FC    SHR   __  __ __       No  16 Gbps - DS8KSU1 5/4/1 WAS 48
_ 2A    198   FCP   SHR   __  __ __       No  16 Gbps - SAN1 CLOUD-P12
_ 38    121   FC    SHR   __  __ __       No  16 Gbps - DS8KSU4 1/4/0 WAS 58
_ 39    1E5   FC    SHR   __  __ __       No  16 Gbps - DS8KSU4 4/4/0 WAS 5B
_ 3A    199   FCP   SHR   __  __ __       No  16 Gbps - SAN2 CLOUD-P12
_ 3B    19D   FCP   SHR   __  __ __       No  16 Gbps - SANF48B-2
_ 44    124   FC    SPAN  01  01 09       No  fctc switch - san64b-a-p9
_ 45    125   FC    SPAN  02  02 09       No  fctc switch - san64b-b-p9
_ 48    158   FC    SHR   01  01 13       No  fc sw - san64b-a-p19
```

Figure 3-7   Channel Path List: Set Filter

4. In the next panel select, channel path type **CIB**, as shown in Figure 3-8.

```
*----------------------- Filter Channel Path List ------------------------*
|                                                                        |
|                                                                        |
| Specify or revise the following filter criteria.                       |
|                                                                        |
| Channel path type  . CIB                                               |
| 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 CIB for channel path type

5. In the next panel, only channel definitions for channel path type CIB are shown. Delete these definitions, as shown in Figure 3-9. Then, 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 . . . . : MUSCB        Lepus
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
d F6    07/1  CIB   SPAN  __  __ __    N   No  Wrap to F7
d F7    07/2  CIB   SPAN  __  __ __    N   No  Wrap to F6
```

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

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

```
*-------------------- Confirm Delete Channel Path ---------------------*
|                                                        Row 1 of 2 | --
| 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 . . . . : MUSCB       Lepus                            |
| Channel Subsystem ID : 0                                            |
|                                                                     |
|                    CHID                                             |
| CHPID   Type  Mode  AID/P                                           |
| F6      CIB   SPAN  07/1                                            |
| F7      CIB   SPAN  07/2                                            |
| ************************* Bottom of data ************************** | **
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
|                                                                     |
*-----------------------------------------------------------------------*
| Spanned channel paths are deleted from all accessing channel subsystems. |
*-----------------------------------------------------------------------*
```

Figure 3-10 Confirm Delete Channel Path panel

3.2.6 Changing the M/T 2965 to 3907 and deleting the 2965

You can keep the original copy of the 2965 (LEPUS) or delete it from the IODF. In this example, keep it in the IODF for a few more steps.

To change the 2965 to a 3907, complete the following steps:

1. Enter `c` (for change) next to MUSCB to change the 2965 to a 3907 and press Enter. The Change Process Definition panel opens (see Figure 3-11 on page 41).
2. Make the following updates and press Enter:
   - Update `Processor type` to 3907.
   - Update `Processor model` to `ZR1`.
   - Update the 2965 part of the `Serial number` to `3907` (that is, `0BB4B7`**`2965`** to `0BB4B7`**`3907`**).
   - Update `Description` to MUSCB.
   - Update `Local system name` to MUSCB.

Note: Because we use HCD to write an IOCDS to this 2965 in preparation for an upgrade, we must leave the *Network name* and *CPC name* set to **IBM390PS** and **LEPUS**. These settings must be updated in the IODF after the 2965 is upgraded to a 3907.

```
*------------------------ Change Processor Definition ------------------------*
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Processor ID . . . . . . . . : MUSCB                                         |
| Support level:                                                               |
| 2965 support, ISM, RCE, UID, CL5                                             |
| Processor type . . . . . . . . 3907      +                                   |
| Processor model  . . . . . . . ZR1       +                                   |
| Configuration mode . . . . . . LPAR      +                                   |
|                                                                              |
| Serial number  . . . . . . . . 0BB4B73907  +                                 |
| Description  . . . . . . . . . MUSCB                                         |
|                                                                              |
| Specify SNA address only if part of an System z cluster:                     |
|                                                                              |
| Network name . . . . . . . . . IBM390PS  +                                   |
| CPC name . . . . . . . . . . . LEPUS     +                                   |
|                                                                              |
| Local system name  . . . . . . MUSCB                                         |
|                                                                              |
|                                                                              |
*------------------------------------------------------------------------------*
```

Figure 3-11 Processors: Change Processor Definition panel

3. The Update Channel Path Identifiers panel opens (see Figure 3-12). No changes are made in this example.

```
*--------------------- Update Channel Path Identifiers ----------------------*
|                                                                Row 1 of 44 |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Specify any changes to the channel path identifiers in the list below.     |
|                                                                            |
| Processor ID . . . . : MUSCB        MUSCB                                  |
| Channel Subsystem ID : 0                                                   |
|                                                                            |
| CHPID  Type  Side  Until CHPID   New CHPID +                               |
| 11     IQD         __            11                                        |
| 28     FC          __            28                                        |
| 29     FC          __            29                                        |
| 2A     FCP         __            2A                                        |
| 38     FC          __            38                                        |
| 39     FC          __            39                                        |
| 3A     FCP         __            3A                                        |
| 3B     FCP         __            3B                                        |
| 44     FC          __            44                                        |
| 45     FC          __            45                                        |
| 48     FC          __            48                                        |
| 49     FC          __            49                                        |
| 4A     FC          __            4A                                        |
*----------------------------------------------------------------------------*
```

Figure 3-12 Processors: Update Channel Path Identifiers panel

4. Press Enter for each Channel Subsystem ID.

The repeated 2965 processor is successfully changed to a 3907-ZR1, as shown in Figure 3-13.

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
_ MUSCB    3907     ZR1      LPAR  0BB4B73907 MUSCB
```

Figure 3-13 Processor List: Changed processor

3.2.7 Deleting the 2965 processor definition

Now that the 2965 was repeated and changed to a 3907, the original 2965 definition (LEPUS) must be deleted so that the required CF Links can be restored.

To delete the 2965 processor definition, complete the following steps:

1. Enter d (for delete) next to the LEPUS processor in the Processor List (see Figure 3-14).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
d LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
_ MUSCB    3907     ZR1      LPAR  0BB4B73907 MUSCB
```

Figure 3-14 Processor List: Deleting processor

2. Press Enter to confirm the deletion of the processor (see Figure 3-15).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ MUSCA    3907     ZR1      LPAR  007A883907 Musca
_ MUSCB    3907     ZR1      LPAR  0BB4B73907 MUSCB
```

Figure 3-15 Processor List: Processor deleted

3.2.8 Reconnecting the CF channel paths that were not migrated

Manually redefine the CF Links that you want from the MUSCB (previously LEPUS) processor to any other processor, along with any Internal Coupling Facility links that you want. To help in this effort, you can get a CF connection report from the previous production IODF containing the 2965. Alternatively, you can make a note of all CBDG441I error messages that you received, as described in 3.2.7, "Deleting the 2965 processor definition" on page 42.

3.3 Saving and restoring OSA configuration data

The three processes for Open Systems Adapter (OSA) cards that you might need to use when upgrading or replacing your processor are described in this section.

3.3.1 Using OSA/SF to save and restore OSE OAT configuration data

For more information about how to save and restore any OSA configuration data such the OSA Address Table (OAT), see 7.3.1, "Saving and restoring OSA-ICC configuration" on page 150.

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

For more information about exporting and importing process for OSA-ICC Server and Session configuration data, see 7.3, "Defining a new OSA-ICC configuration by using OSA Advanced facilities" on page 145.

3.3.3 Using OSA Advanced Facilities to set OSA parameters

For more information about the process of changing the OSA port speed or MAC addresses, see 6.3, "Customizing OSA-Express using OSA Advanced facilities" on page 127.

3.4 Installing a new z14 ZR1 into an existing environment

In this section, we describe the steps to add an IBM z14 ZR1 into an existing IBM Z environment.

3.4.1 Scenario overview

This scenario shows configuration steps for defining a new 3907 processor into an existing hardware environment. This process includes the following key considerations:

- HCD requires a new processor ID for the 3907.
- HCD requires a new CPC name for the 3907.
- The 3907 processor connects to new switch ports and new control unit interfaces.
- The control unit interfaces connect to the same switch ports as they did previously.
- The starting IODF is the current 3906/2965 production IODF.
- The target IODF is a new work IODF with a 3907 defined.

- HCD actions:
  - Migrate updated IOCP statements.
  - Build production IODF.
  - Remote write IODF to IOCDS.
- The HMC actions:
  - Build Reset Profile and point to required IOCDS.
  - Build/verify Image Profiles.
  - Build/verify Load Profiles.
  - Run a power-on reset.

This example defines a new 3907-ZR1 with a Processor ID of MUSCA and with three CSSs (CSS ID=0 to CSS ID=2).

The CPC name of MUSCA and serial number of 02-07A88 are used for the 3907.

The tool requirements are listed in Table 3-2.

Table 3-2 I/O configuration for a new (additional) 3907

New (additional) 3907	**New (additional) 3907 to connect to the new switch ports and same control units 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
Control Unit to switch port connections	Same ports
Starting IODF	Current active production IODF
Target IODF	Create a work IODF
HCD action	Add processor
CHPID Mapping Tool Program	Optional, but good for verifying configuration
`CFReport File` (CCN)	Required for CMT
IOCP (import from validated work IODF)	Yes
CHPID Mapping Tool actions (PCHID reset)	Yes
CHPID Mapping Tool IOCP Output	Yes
CHPID Mapping Tool Reports	Yes, CHIPID Report and CHIPID to CU Report

HCD: Creating a 3907 IODF

The following steps explain how to define an additional 3907 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 new 3907 processor.
3. Using OSA/SF to save and restore OSE OAT configuration data.
4. Export and import OSA-ICC configuration data with OSA Advanced Facilities.
5. Using OSA Advanced Facilities to set OSA parameters.

3.4.2 Creating a work IODF from the current production IODF

HCD is the tool that is used to make a work IODF. However, we start from the current production IODF that contains the existing hardware environment that is connected to the new 3907 (for example, `SYS6.IODF24`).

3.4.3 Adding the new 3907 processor

To add the new 3907 processor, complete the following steps:

1. From the HCD main menu, select option **1.3, Processor List**.
2. In the Processor List (Figure 3-16), press PF11, or enter `add` on the command line to add a processor and press Enter.

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
```

Figure 3-16 Processor List: Adding a processor

The Add Processor panel opens (see Figure 3-17).

```
+------------------------------- 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 System z cluster:                     |
|                                                                             |
| Network name . . . . . . . . . . ________   +                               |
| CPC name . . . . . . . . . . . . ________   +                               |
|                                                                             |
| Local system name  . . . . . . . ________                                   |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap        |
| F12=Cancel                                                                  |
+-----------------------------------------------------------------------------+
```

Figure 3-17 Add Processor: Data fields to be updated

3. Specify the appropriate values. For example, specify the following values, as shown in Figure 3-18:
   - Processor ID: `MUSCA`
   - Processor type: `3907`
   - Processor model: `ZR1`
   - Number of channel subsystems: Blank for now
   - Serial number: `007A883907`
   - Network name: `IBM390PS`
   - CPC name: `MUSCA`
   - Local System Name: Blank for now

```
*------------------------------- Add Processor -------------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . . MUSCA                                      |
| Processor type . . . . . . . . . 3907      +                                |
| Processor model  . . . . . . . . ZR1       +                                |
| Configuration mode . . . . . . . LPAR      +                                |
| Number of channel subsystems . . _         +                                |
|                                                                             |
| Serial number  . . . . . . . . . 007A883907                                 |
| Description  . . . . . . . . . . ________________________________           |
|                                                                             |
| Specify SNA address only if part of a System z cluster:                     |
|                                                                             |
| Network name . . . . . . . . . . IBM390PS  +                                |
| CPC name . . . . . . . . . . . . MUSCA___  +                                |
|                                                                             |
| Local system name  . . . . . . . ________                                   |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 3-18 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, `SYS6.IODF79.WORK`).

5. Press Enter. You now have a 3907 processor named MUSCA (see Figure 3-19).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
_ MUSCA    3907     ZR1      LPAR  007A883907 ______________________________
****************************** Bottom of data ****************************






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

Figure 3-19 Processor List: New processor added

The message at the bottom of the panel indicates that the processor definition is extended to its maximum configuration. This message is generated because part of the main storage is allocated as a fixed-size Hardware Systems Area, which is not addressable by application programs. In HCD, when you define as new or redefine a processor as a 3907, HCD automatically defines the maximum configuration of three CSSs and 40 logical partitions.

6. Enter `s` next to MUSCA and press Enter. The Channel Subsystem List is displayed. Here you can see three channel subsystems (CSS0-CSS2) that are defined with the default MAXDEV values for SS0 of 65280 set by HCD and 65535 set for SS1, and SS2, (see Figure 3-20).

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

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

Processor ID . . . : MUSCA

  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       0         0
_ 1   65280     0       65535     0       65535     0       0         0
_ 2   65280     0       65535     0       65535     0       0         0
```

Figure 3-20 Channel Subsystem List: Three subchannel sets

Tip: In addition to IPL from subchannel set 0 (SS0), IBM Z server with driver level 93 and above support IPL from subchannel set 1 (SS1), subchannel set 2 (SS2), or subchannel set 3 (SS3). Devices that are used early during IPL processing can now be accessed by using subchannel set 1, subchannel set 2, or subchannel set 3.

This configuration allows the users of Metro Mirror (PPRC) secondary devices that are defined by using the same device number and a new device type (3390D) in an alternative subchannel set to be used for IPL, IODF, and stand-alone memory dump volumes when needed.

IPL from an alternative subchannel set is supported by z/OS V1.13 or later, and V1.12 and V1.11 with program temporary fixes (PTFs). IPL applies to the Fibre Channel connection (FICON) and High Performance FICON for z Systems (zHPF) protocols.

3.5 Saving and restoring any OSA configuration data

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

3.5.1 Using OSA/SF to save and restore OSE OAT configuration data

For more information about how to save and restore any OSA configuration data, such as the OAT, see 7.3.1, “Saving and restoring OSA-ICC configuration” on page 150.

3.5.2 Export and import OSA-ICC configuration data with OSA Advanced Facilities

For more information about the exporting and importing process for OSA-ICC Server and Session configuration data, see 7.3, “Defining a new OSA-ICC configuration by using OSA Advanced facilities” on page 145.

3.5.3 Using OSA Advanced Facilities to set OSA parameters

For more information about the process of changing the OSA port speed or MAC addresses, see 6.3.2, “Setting OSA parameters using OSA Advanced Facilities” on page 136.

3.6 Extra steps and processes

When you are ready to map the PCHIDs from the z14 ZR1 CFR file to the CHPIDs in your exported IODF, see Chapter 4, “Mapping CHIDs to CHPIDs by using the CMT” on page 49.

For more information about defining the I/O configuration for your system, see Chapter 5, “Production IODF and setting up the CPC” on page 85.

4

Mapping CHIDs to CHPIDs by using the CMT

This chapter describes a detailed scenario for the use of the CHPID Mapping Tool (CMT).

This chapter includes the following topics:

4.1 Validating the 3907 work IODF

To validate the 3907 work input/output definition file (IODF) by using the hardware configuration definition (HCD), 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 4-1).

```
+---------------- View I/O Definition File Information -----------------+
|                                                                       |
|                                                                       |
| IODF name  . . . . . . : 'ITSO1.IODF78.WORK'                          |
| IODF type  . . . . . . : Work - Validated                             |
| IODF version . . . . . : 5                                            |
|                                                                       |
| Creation date  . . . . : 2018-03-06                                   |
| Last update  . . . . . : 2018-03-08  10:43                            |
|                                                                       |
| Volume serial number . : DZ3P02                                       |
| Allocated space  . . . : 5000    (Number of 4K blocks)                |
| Used space . . . . . . : 2158    (Number of 4K blocks)                |
|    thereof utilized (%)  86                                           |
| Activity logging . . . : No                                           |
| Multi-user access  . . : No                                           |
| Backup IODF name . . . :                                              |
|                                                                       |
| Description  . . . . . :                                              |
|                                                                       |
|                                                                       |
|                                                                       |
|  F1=Help    F2=Split   F3=Exit    F9=Swap   F12=Cancel                |
+-----------------------------------------------------------------------+
```

Figure 4-1 View I/O Definition File Information: Validated work IODF

4.2 Creating the IOCP for the CHPID Mapping Tool

To create the input/output configuration program (IOCP) for the CHPID Mapping Tool, complete the following steps:

1. Select HCD option **2.3. Build IOCP input data set** and press Enter (see Figure 4-2).

```
                      z/OS V2.3 HCD
+------ 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.  *Activate switch configuration             |
|     9.  *Save switch configuration                 |
|     10. Build I/O configuration data               |
|     11. Build and manage System z cluster IOCDSs,  |
|         IPL attributes and dynamic I/O changes     |
|     12. Build validated work I/O definition file   |
|                                                    |
| * = requires TSA I/O Operations                    |
|  F1=Help     F2=Split    F3=Exit     F9=Swap       |
| F12=Cancel                                         |
+----------------------------------------------------+
```

Figure 4-2 Activate or Process Configuration Data: Building IOCP for MUSCA

2. HCD displays the list of available processors (see Figure 4-3). Select the MUSCA processor by entering a forward slash (/) next to it and press Enter.

```
+-------------------------- Available Processors ---------------------------+
|                                                                Row 1 of 3 |
| Command ===> __________________________________________________________   |
|                                                                           |
| Select one.                                                               |
|                                                                           |
|   Processor ID  Type     Model    Mode  Description                       |
|   CETUS         3906     M03      LPAR  Cetus                             |
|   LEPUS         2965     N20      LPAR  Lepus                             |
| / MUSCA         3907     ZR1      LPAR  Musca                             |
| *************************** Bottom of data ****************************   |
|                                                                           |
|                                                                           |
|                                                                           |
|  F1=Help       F2=Split      F3=Exit       F7=Backward   F8=Forward       |
|  F9=Swap      F12=Cancel    F22=Command                                   |
+---------------------------------------------------------------------------+
```

Figure 4-3 Available Processors: Selecting a processor for IOCP file

3. HCD displays a panel on which you enter information about the IOCP input data set to be created (see Figure 4-4). Complete the following fields:
   - Title1: `IODF78`
   - IOCP input data set: `'ITSO1.IODF78.IOCPIN.MUSCA'`
   - 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 . . . . . . . . . : 'ITSO1.IODF78.WORK'                           |
| Processor ID  . . . . . . . : MUSCA                                         |
| Title1 . IODF78                                                             |
| Title2 : ITSO1.IODF78.WORK - 2018-03-08 10:43                               |
|                                                                             |
| IOCP input data set                                                         |
| 'ITSO1.IODF78.IOCPIN.MUSCA'                                                 |
| Input to Stand-alone IOCP?  Yes  (Yes or No)                                |
|                                                                             |
| Job statement information                                                   |
| //WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=T                                   |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
| //*                                                                         |
|                                                                             |
|  F1=Help     F2=Split    F3=Exit     F5=Reset    F9=Swap    F12=Cancel      |
+-----------------------------------------------------------------------------+
```

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

4. Press Enter. HCD submits a batch job to create the data set.

5. In Time Sharing Option (TSO), verify that the data set that you created exists and contains IOCP statements (see Figure 4-5). This data set is used as input into the CHPID Mapping Tool.

```
ID    MSG1='IODF78',                                                  *
      MSG2='ITSO1.IODF78.WORK - 2018-03-08 10:43',                    *
      SYSTEM=(3907,1),LSYSTEM=MUSCA,                                  *
      TOK=('MUSCA',008001117A883907104328120118067F00000000,00*
      000000,'18-03-08','10:43:28','........','........')
RESOURCE PARTITION=((CSS(0),(LEPUS0A,A),(LEPUS0B,B),(LEPUS0C,C*
      ),(LEPUS0D,D),(LEPUS0E,E),(LEPUS0F,F),(LEPUS01,1),(LEPUS*
      02,2),(LEPUS03,3),(LEPUS04,4),(LEPUS05,5),(LEPUS06,6),(L*
      EPUS07,7),(LEPUS08,8),(LEPUS09,9)),(CSS(1),(MUSCA1C,C),(*
      MUSCA1F,F),(MUSCA11,1),(MUSCA12,2),(MUSCA13,3),(MUSCA14,*
      4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,D),(*,E)*
      ),(CSS(2),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8*
      ),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)))
FUNCTION FID=A0,VF=1,PCHID=100,PNETID=PERFNET,                        *
      PART=((MUSCA11),(=)),TYPE=ROC2,PORT=1
FUNCTION FID=A1,VF=2,PCHID=100,PNETID=PERFNET,                        *
```

Figure 4-5 IOCP input data set: Contents (truncated)

Part of the TOK statement is now replaced with dots (see Example 4-1).

Example 4-1 IOCP file (TOK statement)

```
TOK=('MUSCA',008001117A883907104328120118067F00000000,00*
000000,'18-03-08','10:43:28','........','........')
```

These dots ensure that this IOCP file cannot be written to a processor and used for a power-on reset. 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 power-on reset can be generated from a production IODF only.

> **Important:** When an IOCP statement file is exported from a validated work IODF by using HCD, it must be imported back to HCD for the process to be valid. The IOCP file cannot be used directly by the IOCP program.

6. Download this IOCP file from TSO to the CMT workstation. Use a workstation file transfer facility, such as 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 `MUSCAin.iocp`.

4.3 Assigning CHIDs to CHPIDs using the CMT

The following steps use the IOCP statements from HCD steps and the 3907 order process (CFReport). Use the CHPID Mapping Tool to assign CHIDs to each of the CHPIDs for the 3907.

For this process, the CHPID Mapping Tool (CMT) must be downloaded. For more information about downloading and installing the CMT, see 2.5, “CHPID Mapping Tool”. If CMT is installed, verify that the latest updates are installed.

The version of CHPID Mapping Tool that is used for the following figures is 6.19. Check for the latest version on IBM Resource Link.

For more information, see the *CHPID Mapping Tool User's Guide*, GC28-6984.

Use the CHPID Mapping Tool 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 CS5 CHPIDs.
6. Set the priority for single-path control units and other control units that override the CHPID Mapping Tool default priorities and Automatic Mapping.
7. Resolve the CHPIDs that are not connected to control units.
8. Create the CHPID Mapping Tool reports.
9. Create an updated IOCP statements file for transfer back into the IODF file.

4.4 Importing the CFReport file into the CHPID Mapping Tool

To import the CFReport file into the CHPID Mapping Tool, complete the following steps:

1. Start the CMT on your workstation.
2. CMT asks for a project name and location of the CMT work files. We used `MUSCA_upg` as the project name (see Figure 4-6).

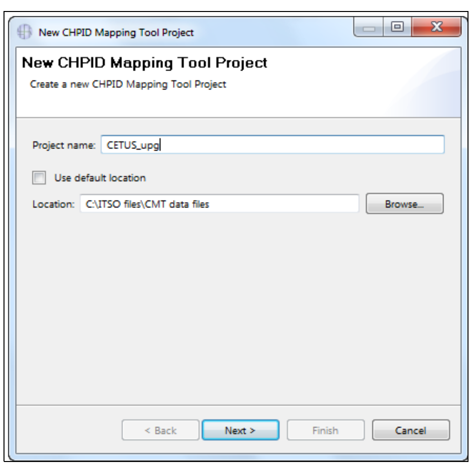

Figure 4-6 Creating a CHPID Mapping Tool Project

3. Specify that the CFReport and IOCP input file panel is displayed. For this step, we input only the CFReport file for now.

> **Attention:** To import the CFReport File into the CHPID Mapping Tool, a Customer Number must be contained in the CFReport File.

4. Import the `CFReport` file into the CHPID Mapping Tool by specifying the name in the CFReport file field. Then, click **Finish** (see Figure 4-7).

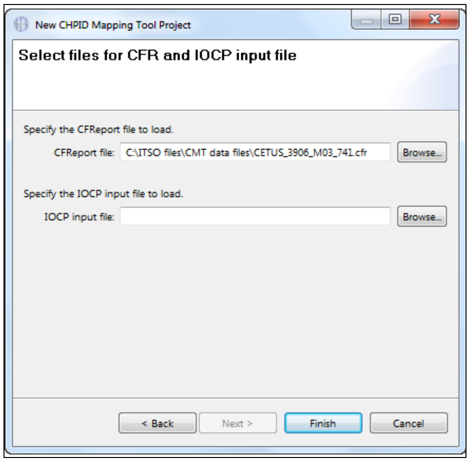

Figure 4-7 Specifying the CFReport file

If you click **Finish** but did not select an IOCP file, you receive the message that is shown in Figure 4-8. Click **OK**.

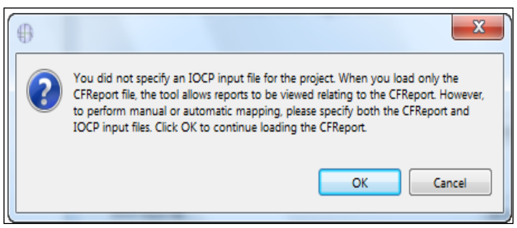

Figure 4-8 Warning message for not specifying an IOCP file

A window shows the progress of reading the CFReport file (see Figure 4-9).

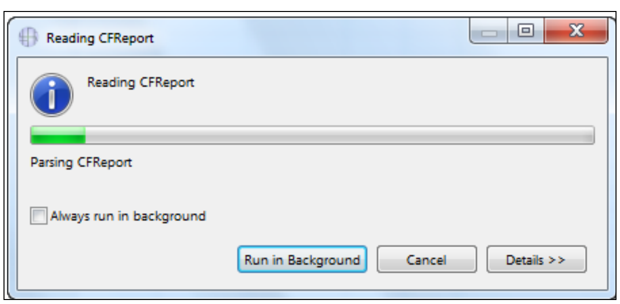

Figure 4-9 Reading the CFReport file

CMT might issue a warning message if it cannot determine that the Adapter IDs in the CFReport file (see Figure 4-10). Click **OK**.

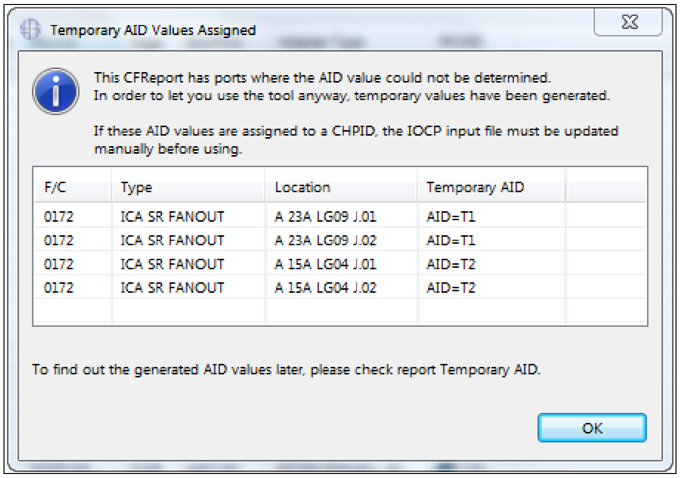

Figure 4-10 Temporary AID Values Assigned

5. The information from the CFReport file is displayed in the Hardware pane (see Figure 4-11).

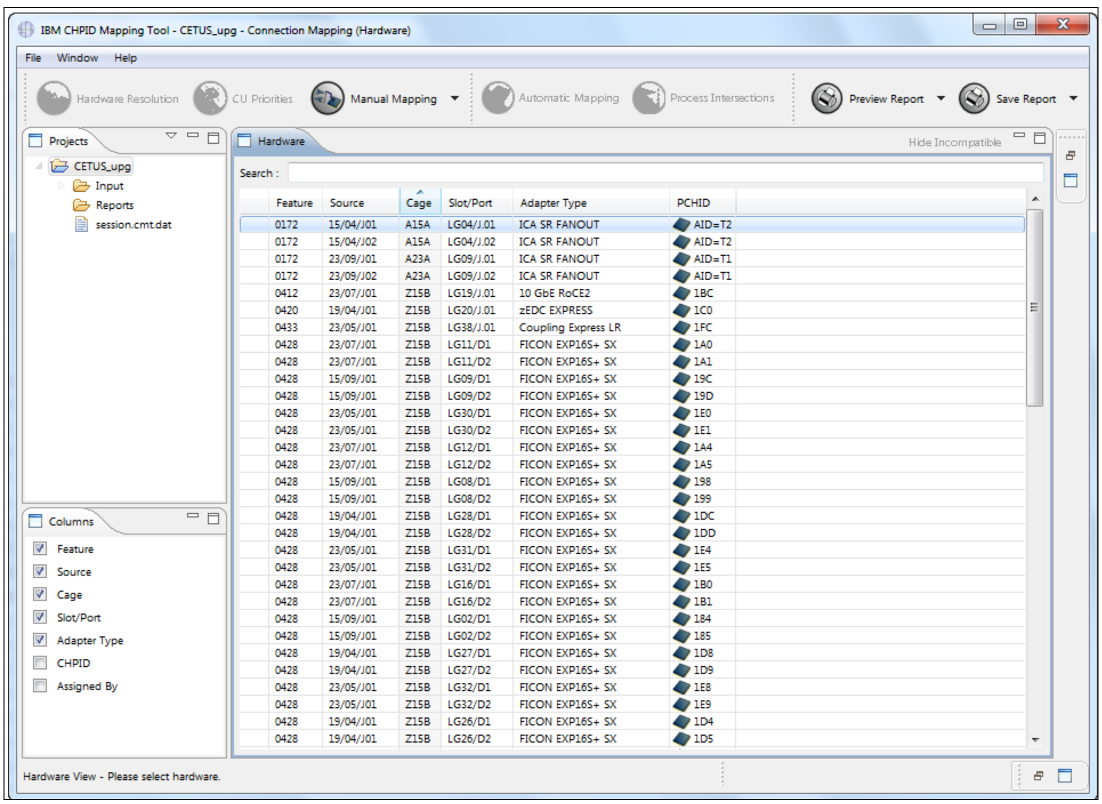

Figure 4-11 Imported CFReport file

4.5 Importing the 3907 IOCP file into the CHPID Mapping Tool

To import the validated 3907 IOCP file into the CHPID Mapping Tool, complete the following steps:

1. Right-click anywhere in the Projects window and select **Import IOCP input file** (see Figure 4-12).

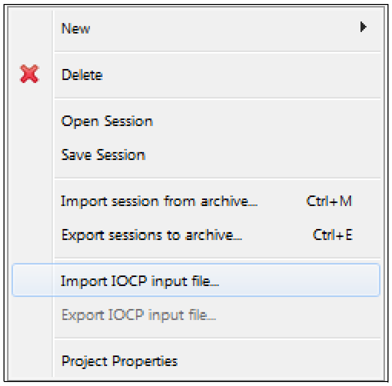

Figure 4-12 Importing the IOCP file

2. Select the IOCP file on your workstation to import into the CHPID Mapping Tool and click **Finish** (see Figure 4-13).

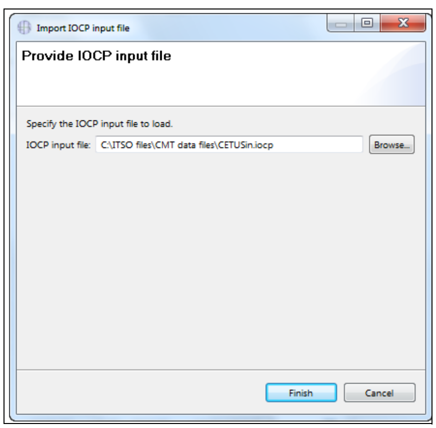

Figure 4-13 Specifying the IOCP file for import

Note: As described in 3.2.5, “Deleting any unsupported items in the repeated 2965” on page 37, z14 ZR1 (machine type 3907) does *not* support InfiniBand coupling links. Therefore, all CHPIDs of type CIB must be deleted in the IOCP file before you import the file into the CHPID Mapping Tool.

3. In the Projects window, under the Input tab, expand the IOCP tab, right-click the IOCP file, and select **Read Selected IOCP** (see Figure 4-14).

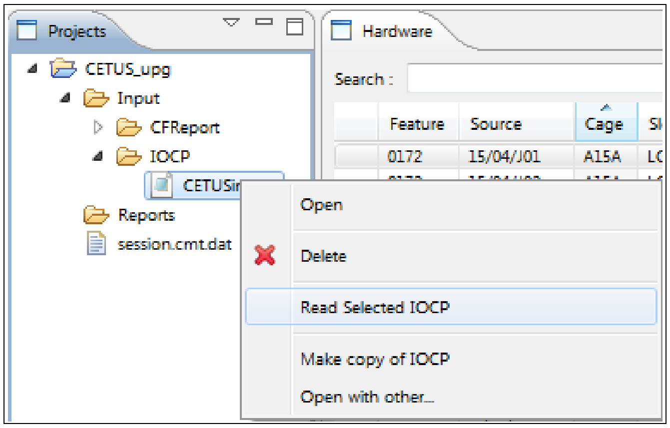

Figure 4-14 Reading the selected IOCP

A window displays the progress information (see Figure 4-15).

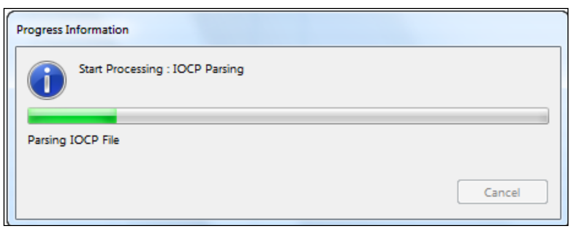

Figure 4-15 Processing the IOCP file

Another window might be displayed that provides a selection regarding what type of upgrade you are performing (see Figure 4-16):

– IOCP file represents current configuration
– IOCP file represents proposed configuration

In our example, we selected **IOCP file represents proposed configuration** because we added I/O during the upgrade process from a 2965 to a 3907. Click **OK**.

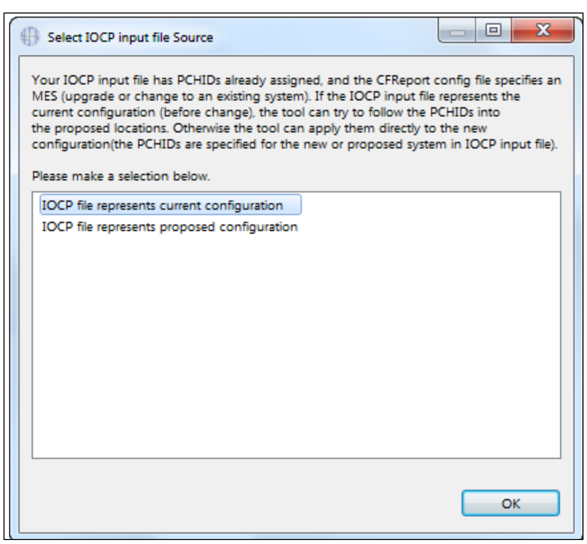

Figure 4-16 Processing the IOCP file

The CHPID Mapping Tool displays the information from the `CFReport` file and the IOCP file in the Hardware Resolution pane. By default, the Hardware Resolution view (see Figure 4-17) includes the following tabbed panes:

– Projects
– Hardware Resolution
– Adapter Type Summary

Hardware Resolution is the middle pane and the Adapter Type Summary is on the right.

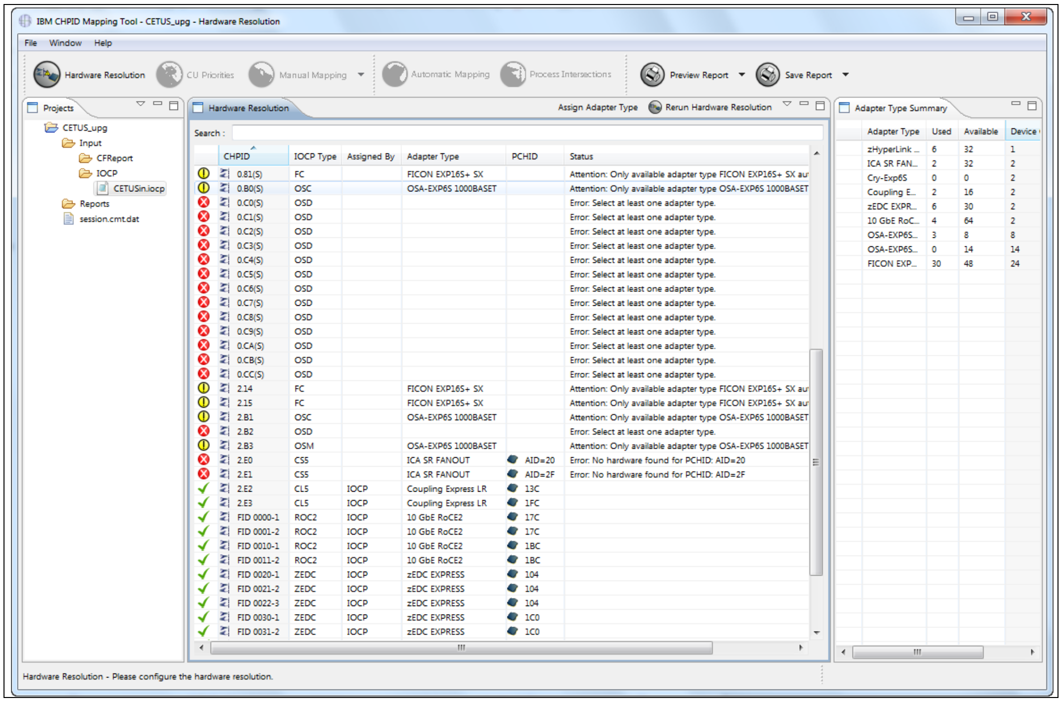

Figure 4-17 Hardware Resolution after Imported IOCPfile

The Adapter Type Summary pane displays a table with helpful information. It summarizes the number of used and available channels for the hardware channel types (used, available, and device count).

In the example, the CHPID Mapping Tool can show some of the following output:

- Hardware Resolution: This window lists all CHPIDs that were found and the Status column shows the CHPID information to be investigated. In the example, investigate the status. The status messages and possible resolutions are listed in Table 4-1.

Table 4-1 Status messages and possible resolutions

Status	Description	Resolution (if required)
No hardware found	AID values or PCHID values are present that are not found in the hardware. This situation might occur when you replace hardware for an MES and the IOCP file contains a CHID value for the old hardware (the IOCP file contains a CHID value for the hardware being removed).	If you use any CHPIDs of IOCP type CS5, the CHPID Mapping Tool cannot automatically assign these CHPIDs. If the AID assignment in the IOCP file is not valid, you can reset it during hardware resolution. You can then use manual mapping to assign the CHPIDs to AIDs. Complete the following steps for CS5 CHPIDs: 1. Remove the AID values. 2. Complete one of the following tasks: – Inside the CHPID Mapping Tool, manually map to associate these CHPIDs with AIDs. – Assign the AID values outside of the tool; for example, by using HCD. 3. Replace the IOCP file.
	Used IOCP types or adapter types exceed the number available. If more IOCP types exist than available, the Adapter Type (in the Hardware Resolution pane) might be empty. If more used than available adapter types exist, the tool indicates this deficiency in the Used and Available columns of the Adapter Type Summary pane.	You must complete one of the following tasks: - Change the IOCP file. - Purchase more hardware. - Ignore the CHPID.
Select at least one adapter type.	An adapter type is not assigned to the current row.	Assign an adapter type to IOCP type.
`Adapter_type` is not compatible with `IOCP_type`.	Adapter type that is assigned for the CHPID is not compatible with the IOCP type that is specified by the IOCP file.	For more information, see Appendix 4.7.1, “Resetting Incompatible (Hardware - I/O) entries” on page 63.
Required hardware for type `IOCP_type` not available. Example: Required hardware for type FC not available.	The CHPID Mapping Tool found no hardware for the specified IOCP type.	You must change the IOCP file or obtain more hardware.
`CHID_1` moved to new channel ID: `CHID_2` **Example**: 520 moved to 1E2	You are replacing hardware for an MES, and the IOCP file contains a CHID value for the old hardware, which is being removed. This CHID value moved from an old machine to the CHID value for the new hardware. `CHID_1` is the first CHID value (for example, 520) and `CHID_2` is the second CHID value (for example, 1E2).	This status is an informational message; no hardware resolution is required. The message informs you of the new location so you can change this location if you prefer a different assignment.

- Manual mapping CS5 CHPIDs: Availability Mapping cannot be used until all CS5 CHPIDs are resolved. You can use manual mapping to resolve any CS5 CHPIDS after which the Availability Mapping function is enabled for use.
- Process the CU Priorities and Automatic Mapping:
  - `Reset CHPIDs assigned by Automatic Mapping`: Selecting this option resets all CHPIDs that were processed by prior availability runs in this session.

    By default, this option is selected.
  - `Reset CHPIDs assigned by Manual Mapping`: Selecting this option resets CHPIDs that were assigned a CHID in the Manual window. If this option is not selected, availability CHIDs for these CHPIDs are not reset.

    By default, this option is not selected.
  - `Reset CHPIDs assigned by IOCP (Potential re-cabling)`: If some of the CHPIDs are assigned in the IOCP Input file, selecting this option resets the CHPIDs. Selecting this option also might require recabling after availability assignments.

    Generally, select this option.
  - `Reset CHPIDs assigned by CMT for config files`: The CFReport indicates that you are performing an MES/upgrade, and channels or CHPIDs (or both) exist that might include configuration files that are associated with them. The MES/upgrade might move some of those channel cards.

    Regardless of whether the channels are moving, the CHPID Mapping Tool assigns CHIDs to the logical CHPID definitions to keep the CHPID definition associated with its current configuration file, or moves the definition to the new location where the channel is moving.

    If you reset the CHPID Mapping Tool assignments, back up the configuration file data before the MES, and restore that data to the new location (the CHID where the affected CHPIDs are assigned) before you use the CHPIDs.

    By default, this option is not selected.

  If no options are selected, availability works on only CHPIDs that do not include assigned CHIDs.

  To give the CHPID Mapping Tool the most choices when you use the availability option, select **Reset CHPIDs assigned by IOCP**.

> **Attention:** If you select **Reset CHPIDs assigned by IOCP**, any mapped CHPID assignments are rest, which can result in recabling of the server.

  However, if you select **Reset CHPIDs assigned by Automatic Mapping**, review the intersects from availability processing carefully to ensure that preserving the previous CHPID-to-CHID relationship does not cause unacceptable availability.

4.6 Resolving CHPIDs with CHID conflicts

The CMT displays the CHPIDs with CHID conflicts (see Figure 4-18).

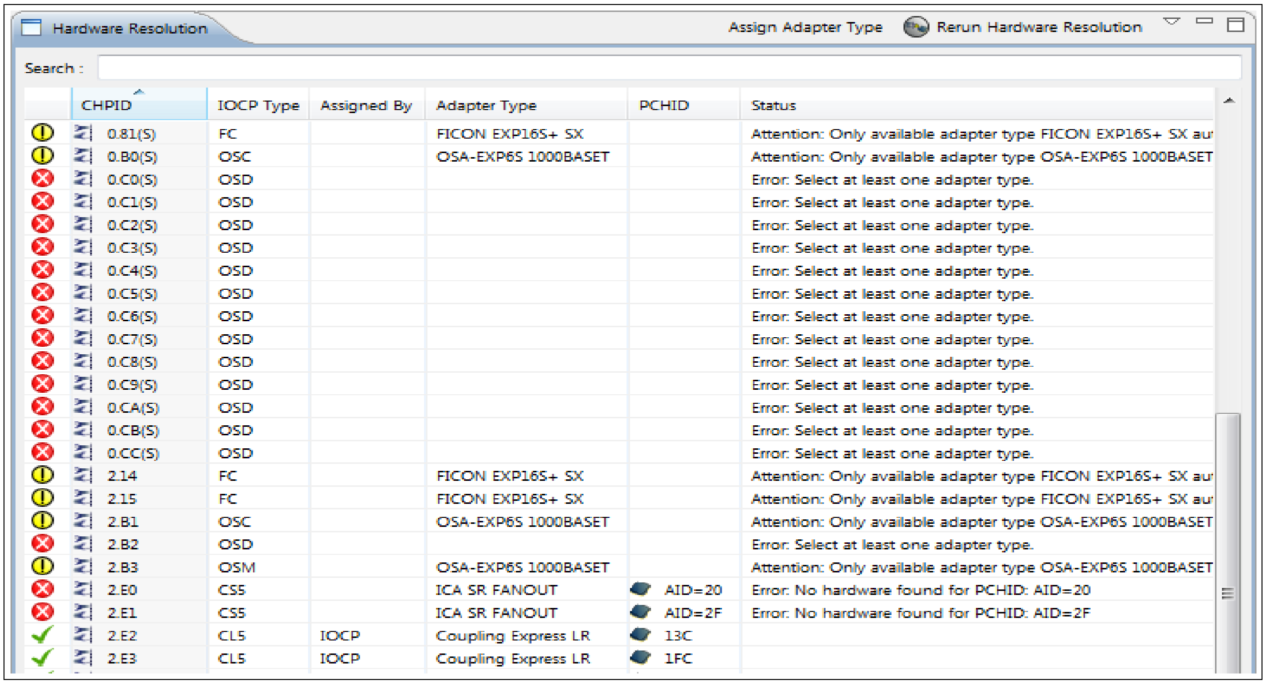

Hardware Resolution | Assign Adapter Type | Rerun Hardware Resolution

Search :

	CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
	0.81(S)	FC		FICON EXP16S+ SX		Attention: Only available adapter type FICON EXP16S+ SX au
	0.B0(S)	OSC		OSA-EXP6S 1000BASET		Attention: Only available adapter type OSA-EXP6S 1000BASET
	0.C0(S)	OSD				Error: Select at least one adapter type.
	0.C1(S)	OSD				Error: Select at least one adapter type.
	0.C2(S)	OSD				Error: Select at least one adapter type.
	0.C3(S)	OSD				Error: Select at least one adapter type.
	0.C4(S)	OSD				Error: Select at least one adapter type.
	0.C5(S)	OSD				Error: Select at least one adapter type.
	0.C6(S)	OSD				Error: Select at least one adapter type.
	0.C7(S)	OSD				Error: Select at least one adapter type.
	0.C8(S)	OSD				Error: Select at least one adapter type.
	0.C9(S)	OSD				Error: Select at least one adapter type.
	0.CA(S)	OSD				Error: Select at least one adapter type.
	0.CB(S)	OSD				Error: Select at least one adapter type.
	0.CC(S)	OSD				Error: Select at least one adapter type.
	2.14	FC		FICON EXP16S+ SX		Attention: Only available adapter type FICON EXP16S+ SX au
	2.15	FC		FICON EXP16S+ SX		Attention: Only available adapter type FICON EXP16S+ SX au
	2.B1	OSC		OSA-EXP6S 1000BASET		Attention: Only available adapter type OSA-EXP6S 1000BASET
	2.B2	OSD				Error: Select at least one adapter type.
	2.B3	OSM		OSA-EXP6S 1000BASET		Attention: Only available adapter type OSA-EXP6S 1000BASET
	2.E0	CS5		ICA SR FANOUT	AID=20	Error: No hardware found for PCHID: AID=20
	2.E1	CS5		ICA SR FANOUT	AID=2F	Error: No hardware found for PCHID: AID=2F
	2.E2	CL5	IOCP	Coupling Express LR	13C	
	2.E3	CL5	IOCP	Coupling Express LR	1FC	

Figure 4-18 CHPIDs with PCHID conflicts

In the first column of every row, the Hardware Resolution pane contains one of the following symbols:

- An X in a red circle: This symbol indicates an error.
- An exclamation mark in a yellow circle: This symbol indicates a warning or attention message.
- A green check mark: This symbol indicates that the tool successfully resolved the specified Channel Type.

The example includes the following reasons to resolve hardware resolution issues:

- The CHID channel type changed.
- The defined CHID is not compatible with the channel path at a particular location.
- Not enough ports exist in the hardware.
- A type mismatch exists between a CHPID and its associated channel type.

4.7 Resolving hardware issues

In the example, the CHPID Mapping Tool displays an X in a red circle in the first column of the Hardware Resolution pane (see Figure 4-19) that is related to error types: `No hardware found` and `FICON EXP16S+ SX is not compatible with OSD`.

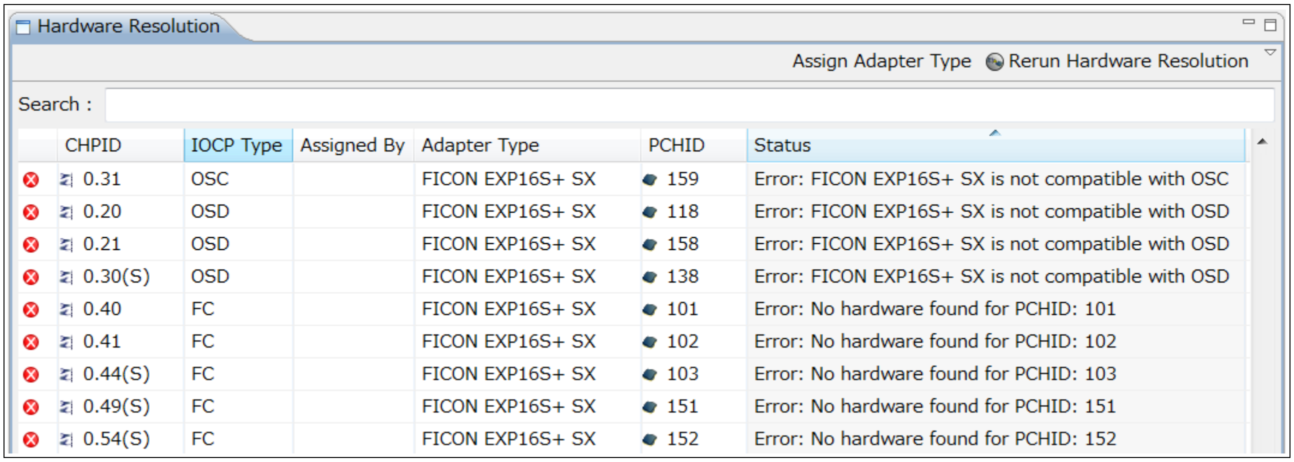

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
0.31	OSC		FICON EXP16S+ SX	159	Error: FICON EXP16S+ SX is not compatible with OSC
0.20	OSD		FICON EXP16S+ SX	118	Error: FICON EXP16S+ SX is not compatible with OSD
0.21	OSD		FICON EXP16S+ SX	158	Error: FICON EXP16S+ SX is not compatible with OSD
0.30(S)	OSD		FICON EXP16S+ SX	138	Error: FICON EXP16S+ SX is not compatible with OSD
0.40	FC		FICON EXP16S+ SX	101	Error: No hardware found for PCHID: 101
0.41	FC		FICON EXP16S+ SX	102	Error: No hardware found for PCHID: 102
0.44(S)	FC		FICON EXP16S+ SX	103	Error: No hardware found for PCHID: 103
0.49(S)	FC		FICON EXP16S+ SX	151	Error: No hardware found for PCHID: 151
0.54(S)	FC		FICON EXP16S+ SX	152	Error: No hardware found for PCHID: 152

Figure 4-19 Hardware resolution status errors

> **More information:** For more information about these error messages, see the *CHPID Mapping Tool User's Guide*, GC28-6947.

The following options must be reset:

- Incompatible (Hardware - I/O) entries
- "Error: No hardware found"
- "Select at least one adapter type"
- "Required hardware for type IOCP_type not available"
- "CHID_1 moved to new channel ID: CHID_2"

4.7.1 Resetting Incompatible (Hardware - I/O) entries

The Channel type that is assigned for the CHPID is not compatible with the IOCP type that is specified by the IOCP file. For this mismatch, you might receive the following message:

`Error: `*`Channel_type`*` is not compatible with `*`IOCP_type`*`.`

Resolve this problem by resetting the CHID. In our example, the IOCP type is OSD, but the CHID is associated with an FICON card. You cannot assign the OSD type on the FICON card.

The CHPID Mapping Tool displays the error message in the Status column, as shown in Figure 4-20.

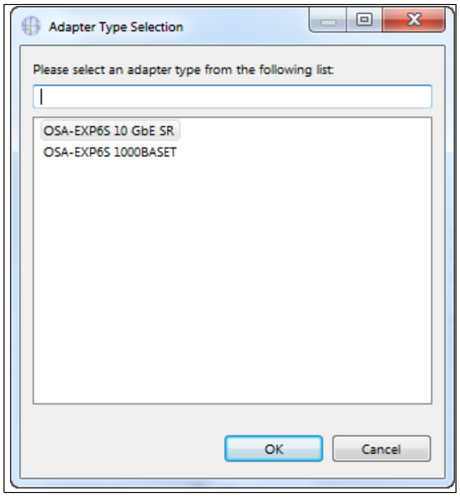

Figure 4-20 Channel_type is not compatible with IOCP_type

Select the channel type **OSD**. The Status is `Error: FICON EXP16S+ SX is not compatible with OSD`. Right-click in the row and select **Reset Incompatible (Hardware - I/O) Entries** to remove the CHID values for only those rows (see Figure 4-21).

Figure 4-21 Channel_Type is not compatible with IOCP_type OSD

The tool replaces the X in a red circle with an Attention icon (exclamation mark in a yellow circle), changes the status message, and removes the CHIDs information (see Figure 4-22).

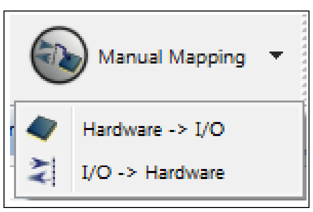

Figure 4-22 Results for reset of incompatible

The CHPID Mapping Tool now displays messages about any CHPID types that were imported from the IODF into the CMT that do not have any associated hardware support in the `CFReport` file (see Figure 4-23). Click **OK**. The Adapter Type Summary details also is shown in Figure 4-23.

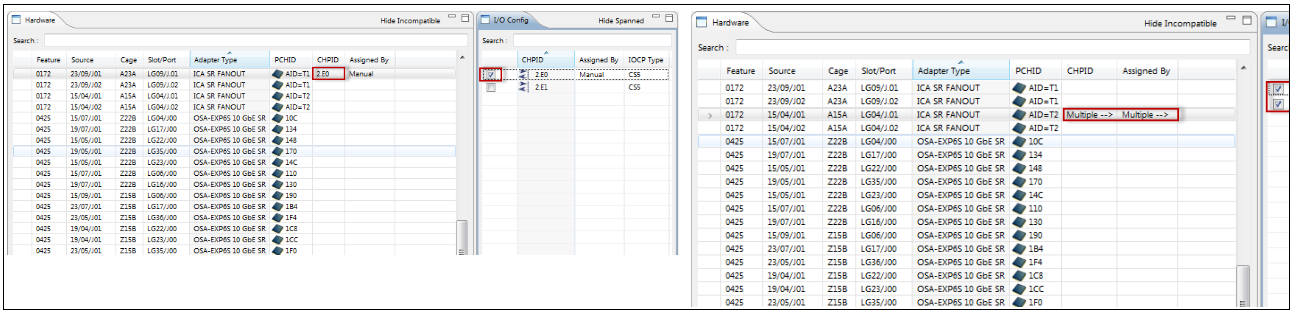

Figure 4-23 Required Hardware unavailable

Excessive numbers of OSD CHPID types are in the example IODF to show how the CHPID Mapping Tool handles this condition.

You can use the overdefine option to change the CHID value to an asterisk (*) in the IODF. In this way, you can retain the OSD CHPID definitions in the IODF so that you can install OSD CHIDs in the processor later.

> **Tip:** Other CHPID types can also be *overdefined* by entering an asterisk (*) for the CHID value. Overdefining is now supported for CS5 type CHPID definitions.

Alternatively, you can remove the OSD CHPID definitions from the IODF.

To continue with our example, complete the following steps:

1. Return to the IODF and change the CHID values for the OSD CHPIDs (or any other CHPIDs that do not include any supporting hardware in the CFReport) to an asterisk (*).
2. Revalidate the IODF by using HCD option 2.12.
3. Re-create the IOCP statements file and transfer it to your workstation.
4. Import the IOCP file by right-clicking the Projects window and selecting **Import IOCP File**.

> **Tip**: If you review the IOCP statements file now, the OSD CHPIDs are still defined in the IODF (although they are omitted from the file).

Now, when you click **Reset "*Channel-Type* is not compatible with IOCP_*type*"**, the CHPID Mapping Tool prompts you to resolve some hardware.

4.7.2 Resetting “Error: No hardware found”

An X in a red circle in the first column indicates an error, and the Status column provides the information with value of `Error: No hardware found` (see Figure 4-24).

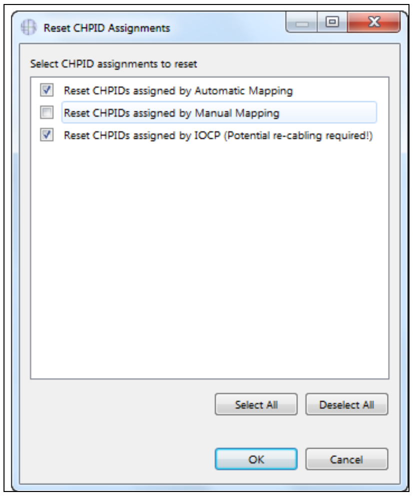

Figure 4-24 Error: No Hardware found

In the example, select channel type **FC**. The Status is `Error: No Hardware found`. Right-click in the row and select **Reset “No hardware found” Entries** to remove the CHID values for those rows (see Figure 4-25).

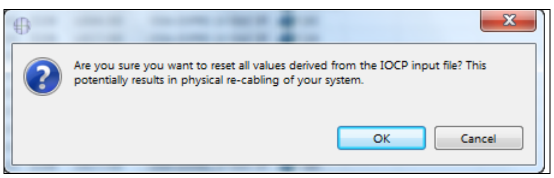

Figure 4-25 Resetting No Hardware found entries

The tool replaces the X with an Attention icon, changes the status message, and removes the CHID information (see Figure 4-26).

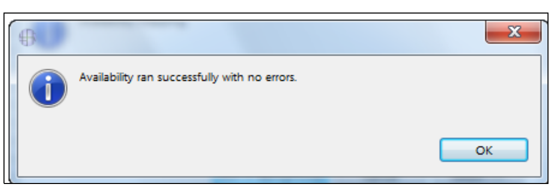

Figure 4-26 Results of resetting No hardware found

4.7.3 Resetting “Select at least one adapter type”

The adapter type is not assigned to the current row. Assign an adapter type to the IOCP type by completing the following steps:

1. Click the Adapter Type column in the target row. The tool displays an arrow in the Channel Type column of the target row (see Figure 4-27).

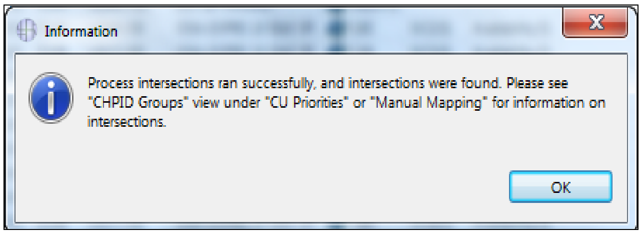

Figure 4-27 Selecting at least one adapter type

2. Click the **ellipses** (...) box.
3. The tool displays a list of available and compatible card types for the CHPID, as shown in Figure 4-28. Select an adapter type and click **OK**.

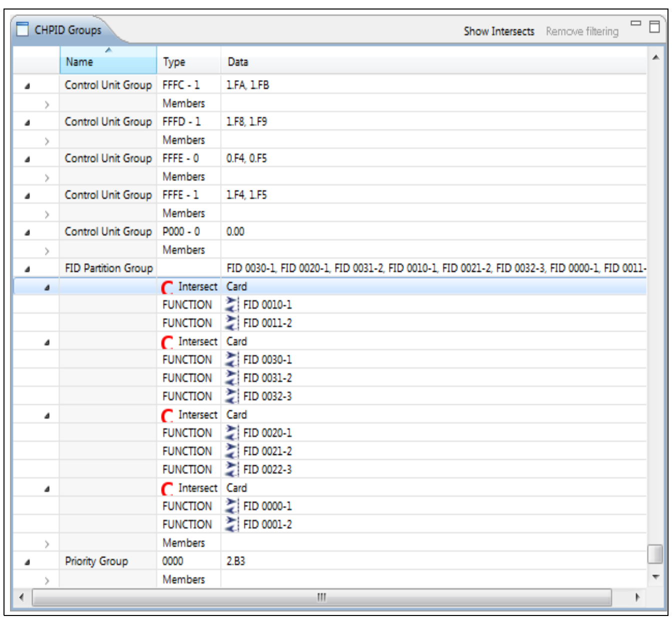

Figure 4-28 Adapter Type Selection

The Used and Available totals change is shown in the Adapter Type Summary tab.

4.7.4 Resetting “Required hardware for type IOCP_type not available”

The CHPID Mapping Tool found no hardware for the specified IOCP type, as shown in the following example:

```
Required hardware for type CS5 not available.
```

You must change IOCP file or obtain more hardware.

4.7.5 Resetting “CHID_1 moved to new channel ID: CHID_2”

When moving from old hardware to new hardware (for example, during a miscellaneous equipment specification [MES]), the CHID value that is assigned to a feature can change. This message indicates that the IOCP file contains a CHID value for the old machine that is being removed. The CHID value is changed from the old machine to the CHID value for the new machine.

For example, CHID_1 is the first CHID value that represents the old hardware (for example, 1B0) and CHID_2 is the new value representing the new hardware (for example, 533). In essence, the feature is present in the old and new hardware, but its location (CHID) changed.

This status is an informational message. No hardware resolution is required. The message informs you of the new location so you can change it if you prefer a different assignment.

After you assign all Adapter Types, the Manual Mapping option becomes available.

4.8 Manual mapping to resolve CS5 CHPIDs

In some situations, the Automatic Mapping option is not available. You cannot use automatic mapping until all CS5 CHPIDs are resolved. You can use manual mapping to resolve this task.

To resolve the CS5 CHPIDs, assign the available CHPIDs by completing the following steps:

1. Click **Manual Mapping** (see Figure 4-29).

Figure 4-29 Manual Mapping

2. Ensure that the tool is set to display Manual Mapping in the Hardware → I/O format (see Figure 4-30).

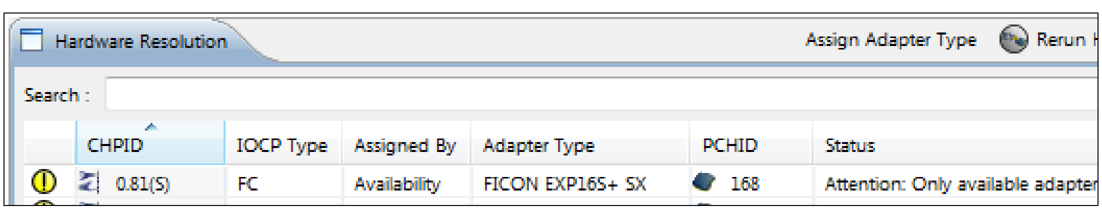

Figure 4-30 Manual Mapping of Hardware → I/O

3. Click every row that includes type ICA SR in the Adapter Type column. The tool displays all the available CHPIDs with IOCP type (see Figure 4-31).

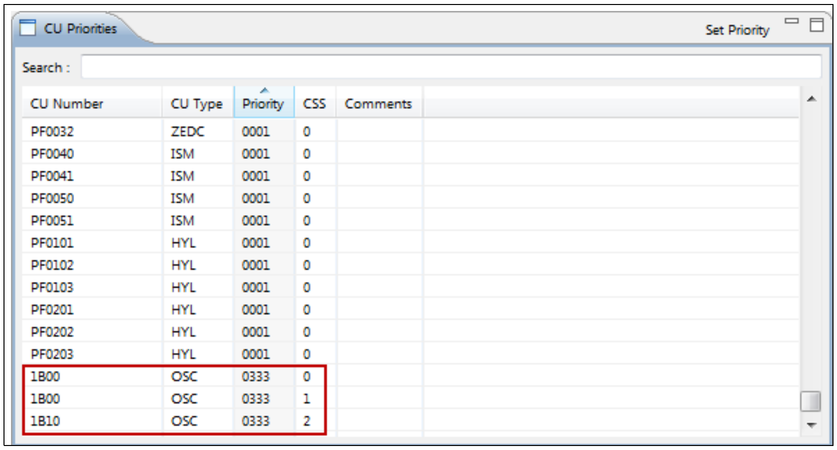

Figure 4-31 Adapter Type of HCA3 and associated CHPID assigned

4. Select one or more empty check boxes in the I/O Config pane to assign the CHPID. In the Hardware pane, the `CHPID number` is inserted in the CHPID column. In the Assigned By column, the value of `Manual` is inserted.

   If you select more than one CHPID for an ICS SR adapter type, you see the `Multiple -->` value (see Figure 4-32) inserted in the CHPID and Assigned By columns.

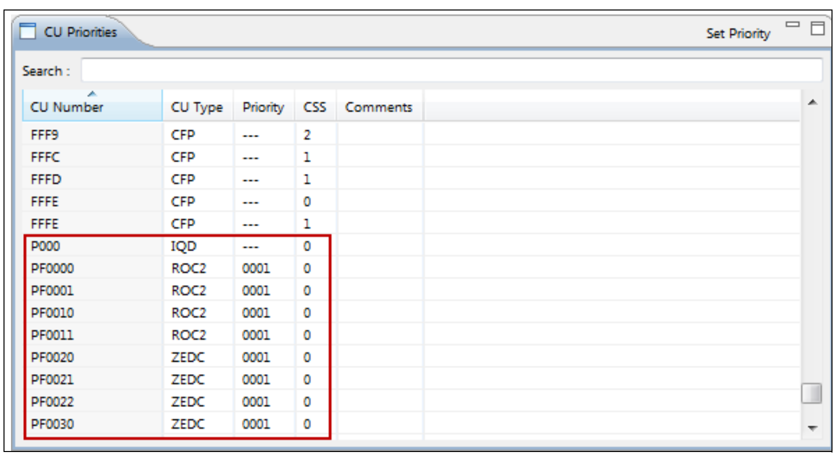

Figure 4-32 Adapter Type of HCA3 and associated multiple CHPID assigned

The Automatic Mapping option becomes available after you assign all the CHPIDs of IOCP type CS5.

4.9 Processing Automatic Mapping then CU Priority

If you are importing an IOCP statements file from a 2965 in which CU Priority values were defined, review the CU Priority values first. The CHPID Mapping Tool can then perform the availability functions for a 3907.

You must assign priorities if you want to make some control units more important (in the CMT processing order) than others, or have two (or more) control units that you want the CMT to process at the same time.

Perform the first availability function by completing the following steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices (see Figure 4-33). For the example, select the following two options and then, click **OK**:
   - **Reset CHPIDs assigned by Automatic Mapping**
   - **Reset CHPIDs assigned by IOCP (Potential re-cabling required!)**

Figure 4-33 Reset CHPID Assignments

> **Tip:** The following fourth choice is also available, but only for an upgrade or an MES:
>
> `Reset CHPIDs assigned by CMT for config files.`

3. Click **OK** to confirm the reset (see Figure 4-34).

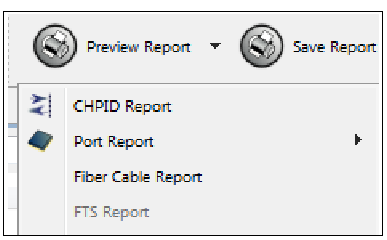

Figure 4-34 Reset CHPID assignments warning message

4. Because the 3907 includes availability rules that differ from a 2965, remove all CHID assignments that are still in the IOCP.
5. Click **OK**.

6. After the CHPID Mapping Tool resets the CHPIDs, it displays the result of the process (see Figure 4-35). Click **OK**.

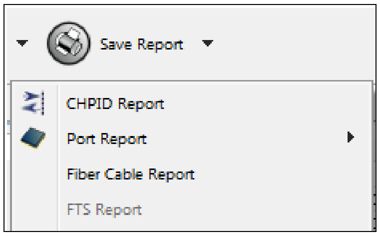

Figure 4-35 Availability ran successfully with no errors message

7. Click **OK** (see Figure 4-36).

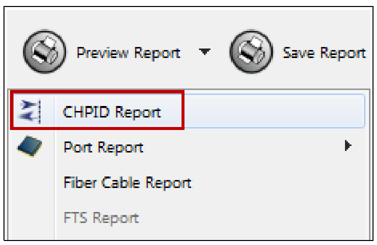

Figure 4-36 Process Intersections run successfully message

The following intersects are available:

C Two or more assigned channels use the same channel card.

S More than half the assigned channels use the same InfiniBand or STI link.

M All assigned channels are supported by the same MBA group.

B More than half the assigned channels are supported by the same MBA Group.

D Assigned channels are on the same daughter card.

> **Tip:** Intersect messages inform you of a potential availability problem that was detected by the CMT. However, they do not necessarily indicate an error. It is your responsibility to evaluate whether the condition must be corrected.

8. Click **Manual Mapping**. In the CHPID Groups tab, observe any intersect warnings that were found during automatic mapping and decide whether they are acceptable (see Figure 4-37). The example returned the “C” intersect. This warning indicates that multiple definitions exist on the same I/O card.

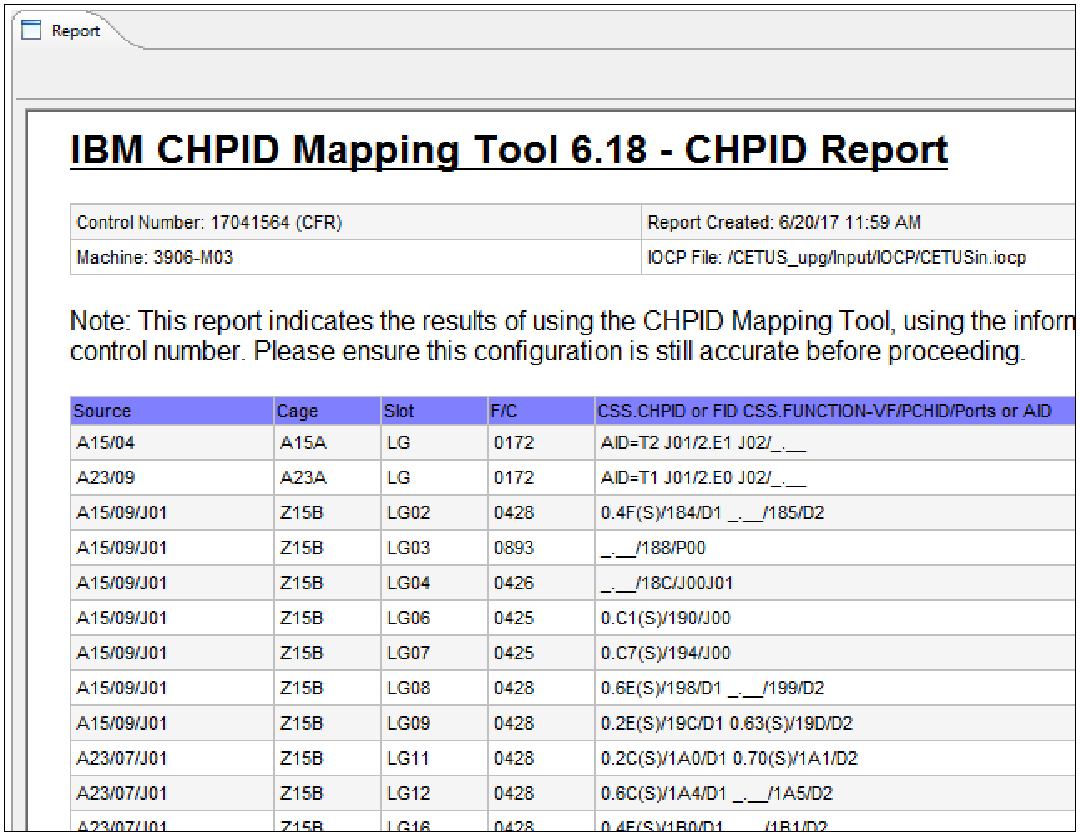

Report

IBM CHPID Mapping Tool 6.18 - CHPID Report

Control Number: 17041564 (CFR)	Report Created: 6/20/17 11:59 AM
Machine: 3906-M03	IOCP File: /CETUS_upg/Input/IOCP/CETUSin.iocp

Note: This report indicates the results of using the CHPID Mapping Tool, using the inforn control number. Please ensure this configuration is still accurate before proceeding.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
A15/04	A15A	LG	0172	AID=T2 J01/2.E1 J02/_.__
A23/09	A23A	LG	0172	AID=T1 J01/2.E0 J02/_.__
A15/09/J01	Z15B	LG02	0428	0.4F(S)/184/D1 _.__/185/D2
A15/09/J01	Z15B	LG03	0893	_.__/188/P00
A15/09/J01	Z15B	LG04	0426	_.__/18C/J00J01
A15/09/J01	Z15B	LG06	0425	0.C1(S)/190/J00
A15/09/J01	Z15B	LG07	0425	0.C7(S)/194/J00
A15/09/J01	Z15B	LG08	0428	0.6E(S)/198/D1 _.__/199/D2
A15/09/J01	Z15B	LG09	0428	0.2E(S)/19C/D1 0.63(S)/19D/D2
A23/07/J01	Z15B	LG11	0428	0.2C(S)/1A0/D1 0.70(S)/1A1/D2
A23/07/J01	Z15B	LG12	0428	0.6C(S)/1A4/D1 _.__/1A5/D2
A23/07/J01	Z15B	LG16	0428	0.4E(S)/1B0/D1 _.__/1B1/D2

Figure 4-37 B Intersect examples

You can now display the results of the channel mapping. You can also sort the report in various ways. For example, you can see how the CHPID Mapping Tool ranked control units.

Complete the following steps to check and set values for items, such as OSC CHPIDs and FCTC CHPIDs, to ensure that the CHPID Mapping Tool allocates these CHPIDs with high CHID availability:

1. Click **CU Priorities**. By default, this pane is in the center at the top of the window.
2. In the CU Priorities pane, search in the CU Number column for the control units for which you want to set a priority.
3. Enter a priority number for the CU in the Priority column for each row. The CHPID Mapping Tool makes more related changes in the CHPID Groups panes.

4.10 CHPIDs not connected to control units

In the CU Priorities window, click in the **CU Number** column (see Figure 4-38). The CHPID Mapping Tool shows at the end of the list all CHPIDs that are defined in the IOCP input that are not connected to control units. In the list of CU numbers, the letter "S" precedes all coupling CHPIDs, and the letter "P" precedes all non-coupling CHPIDs.

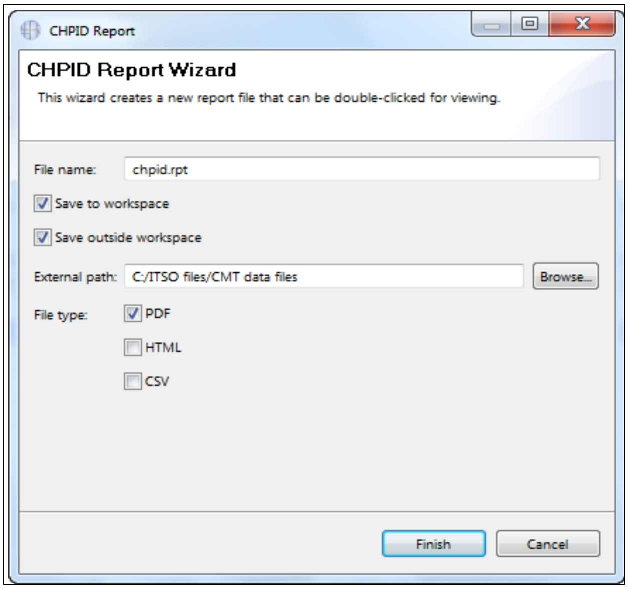

Figure 4-38 CHPIDs not connected to control units

Review the list for the following reasons:

- Perhaps you forgot to add a CHPID to a control unit and must update the IOCP source before you continue in the CMT.
- The unconnected CHPIDs might be extra channels that you are ordering in anticipation of new control units.
- The unconnected CHPIDs might be coupling links that are being used in coupling facility (CF) images (they do not require control units).

If extra CHPIDs exist for anticipated new control units, consider grouping these CHPIDs with a common priority. Having a common priority allows the availability mapping function to pick CHIDs that can afford your new control unit availability.

4.11 Creating CHPID Mapping Tool reports

The CHPID Mapping Tool offers built-in reports, which are available from the top of the window. You can also print the information from the report by clicking **Print**. The options to create a Preview Report or Save Report are listed in Figure 4-39.

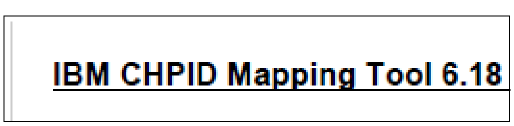

Figure 4-39 Preview Report and Save Report buttons

Click **Preview Report** or **Save Report** to display choices (a list of types of reports). The choices are the same except that **Save Report** lists an extra selection (see Figure 4-40).

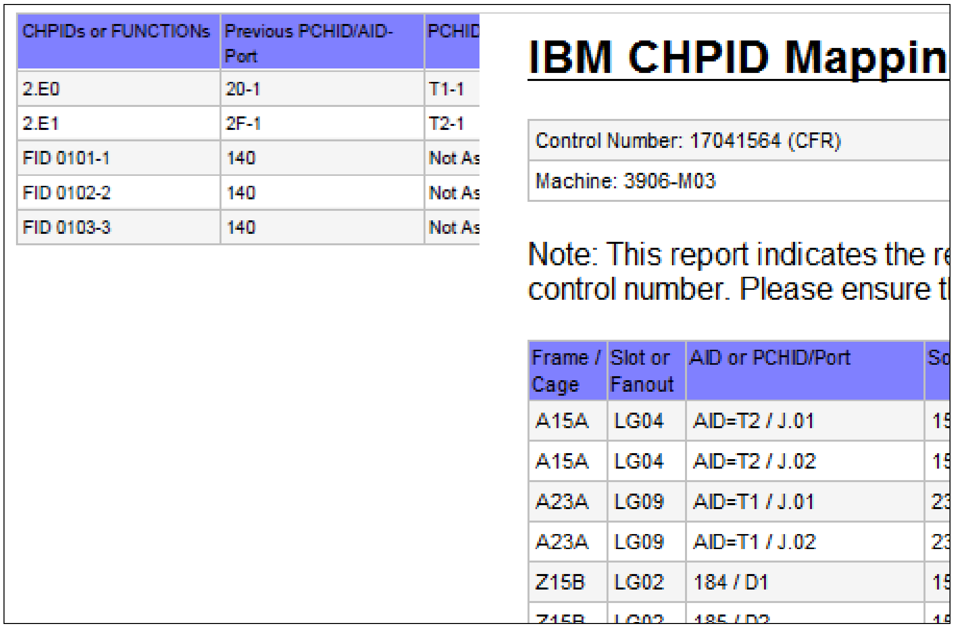

Figure 4-40 Preview Report and Save Report menus

For simplicity, only the following reports are described in this example:

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

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

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/port information before system delivery.

4.11.1 CHPID Report

To create the CHPID report, complete the following steps:

1. Click **Preview Report** → **CHPID Report** (see Figure 4-41).

IBM CHPID Mapping Tool 6.18 - CHPID to CU R

Control Number: 17041564 (CFR)	Report Created: 6/20/17 12:43 PM
Machine: 3906-M03	IOCP File: /CETUS_upg/Input/IOCP/CETUSin.iocp

Note: This report indicates the results of using the CHPID Mapping Tool, using
control number. Please ensure this configuration is still accurate before procee

CSS	CHPID	Type	Source	Port	PCHID/AID-Port
0	00	IQD	N/A	N/A	N/A
0	2C	FC	23/07/J01	Z15B LG11 D1	1A0

Figure 4-41 Preview report: CHPID Report

The CHPID Mapping Tool displays the CHPID Report in a Report tab within the CMT (see Figure 4-42).

1	C6	OSD	23/05/J01	Z15B LG35 J00	1F0	1C60	OSA	---
1	C7	OSD	15/09/J01	Z15B LG07 J00	194	1C70	OSA	---
1	C8	OSD	15/07/J01	Z22B LG06 J00	110	1C80	OSA	---
1	C9	OSD	19/07/J01	Z22B LG16 J00	130	1C90	OSA	---
1	CA	OSD	15/05/J01	Z22B LG23 J00	14C	1CA0	OSA	---
1	CB	OSD	23/07/J01	Z15B LG17 J00	1B4	1CB0	OSA	---
1	CC	OSD	19/05/J01	Z22B LG35 J00	170	1CC0	OSA	---
1	F0	IQD	N/A	N/A	N/A	F000	IQD	---
1	F4	ICP	N/A	N/A	N/A	FFFE	CFP	---
1	F5	ICP	N/A	N/A	N/A	FFFE	CFP	---
1	F8	ICP	N/A	N/A	N/A	FFFD	CFP	---
1	F9	ICP	N/A	N/A	N/A	FFFD	CFP	---
1	FA	ICP	N/A	N/A	N/A	FFFC	CFP	---
1	FB	ICP	N/A	N/A	N/A	FFFC	CFP	---
2	14	FC	19/04/J01	Z15B LG26 D1	1D4	A000	2107	---
						0041	2032	---
2	15	FC	15/07/J01	Z22B LG07 D1	114	A000	2107	---
2	B1	OSC	15/07/J01	Z22B LG03 J00J01	108	1B10	OSC	0333
2	B2	OSD	19/04/J01	Z15B LG22 J00	1C8	1B20	OSA	---
2	B3	OSM	23/07/J01	Z15B LG18 J00J01	1B8	1B30	OSM	0000
2	E0	CS5	23/09/J01	A23A LG09 J.01	AID=T1	FFF9	CFP	---
2	E1	CS5	15/04/J01	A15A LG04 J.01	AID=T2	FFF9	CFP	---
2	E2	CL5	19/07/J01	Z22B LG19 J.01	13C	FFF9	CFP	---
2	E3	CL5	23/05/J01	Z15B LG38 J.01	1FC	FFF9	CFP	---
2	E4	ICP	N/A	N/A	N/A	FFF9	CFP	---
2	E5	ICP	N/A	N/A	N/A	FFF9	CFP	---

Figure 4-42 CHPID Report

Tip: You can save individual reports as multiple reports in batch.

2. Click **Save Report**.

In the example, when you click **CHPID Report**, an option window opens (see Figure 4-43). Specify a file name and an external path (location) of where to save the file. If you want to save the report in HTML, select **HTML**. The tool selects **PDF** by default. The window is similar for all type of reports. Click **Finish**.

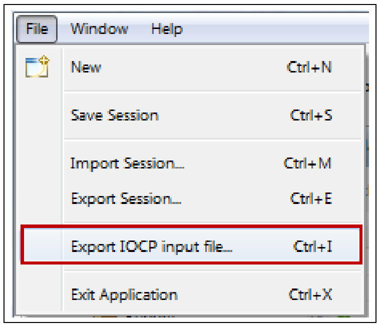

Figure 4-43 Save CHPID Report

The CHPID Report is created by the CHPID Mapping Tool (see Figure 4-44).

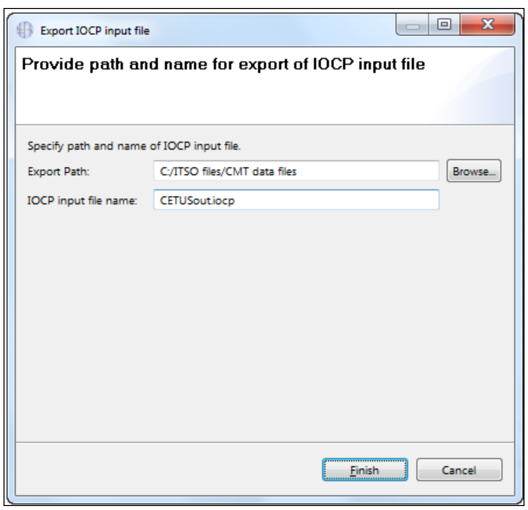

Figure 4-44 CHPID Report example in PDF format

At the end of this CHPID Report is a list of CHPIDs with modified CHID/AID assignments (see Figure 4-45). This report is valuable for moving cables.

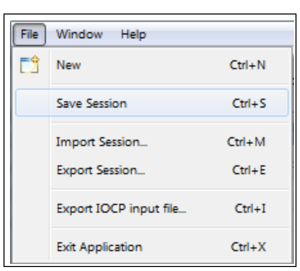

Figure 4-45 List of CHPIDs that include modified PCHID/AID assignments

4.11.2 CHPID to Port Report, sorted by location

To create the Port Report that is sorted by location, click **Preview Report** → **Port Report** → **Sorted by Location**. The CHPID Mapping Tool displays the CHPID to Port Report in a Report tab within the CMT (see Figure 4-46).

Report

Print

IBM CHPID Mapping Tool 6.19 - CHPID to Port Report

Control Number: 02433073 (CFR)	Report Created: 3/29/18 5:00 AM
Machine: 3907-ZR1	IOCP File: /MUSCA_upg/Input/IOCP/MUSCAin.iocp

Note: This report indicates the results of using the CHPID Mapping Tool, using the information based on the above control number. Please ensure this configuration is still accurate before proceeding.

Frame / Cage	Slot or Fanout	AID or PCHID/Port	Source	Adapter Type	Assigned CHPID or Assigned FUNCTION	CHPID Origin or Function Origin
A01B	LG02	140 / J.01	09/04/J01	zEDC EXPRESS	FID D0-1 , FID D1-2	IOCP
A01B	LG03	144 / J.01	09/04/J01	Coupling Express LR	1.E8	Availability
A01B	LG03	144 / J.02	09/04/J01	Coupling Express LR	1.E9	Availability
A01B	LG04	148 / D1	09/04/J01	OSA-EXP6S 10 GbE SR	0.B8(S)	Availability(S)
A01B	LG05	14C / D1D2	09/04/J01	OSA-EXP6S 1000BASET	0.B2(S)	Availability(S)
A01B	LG07	150 / D1D2	09/04/J01	OSA-EXP6S 1000BASET	0.B0(S)	Availability(S)
A01B	LG09	158 / D1	09/04/J01	FICON EXP16S+ SX	0.49(S)	Availability(S)
A01B	LG09	159 / D2	09/04/J01	FICON EXP16S+ SX		
A01B	LG10	15C / D1	09/04/J01	FICON EXP16S+ SX	0.48(S)	Availability(S)
A01B	LG10	15D / D2	09/04/J01	FICON EXP16S+ SX		
A01B	LG12	160 / J.01	09/10/J01	10 GbE RoCE2	FID B0-1 , FID B2-3	IOCP
A01B	LG12	160 / J.02	09/10/J01	10 GbE RoCE2	FID B1-2 , FID B3-4	IOCP
A01B	LG13	164 / P.00	09/10/J01	Cry-Exp6S		
A01B	LG14	168 / D1	09/10/J01	zHyperLink Express	FID 200-1 , FID 202-3	IOCP
A01B	LG14	168 / D2	09/10/J01	zHyperLink Express	FID 201-2 , FID 203-4	IOCP
A01B	LG15	16C / D1	09/10/J01	OSA-EXP6S 10 GbE SR	0.B9(S)	Availability(S)
A01B	LG17	170 / D1D2	09/10/J01	OSA-EXP6S 1000BASET	0.B4(S)	Availability(S)
A01B	LG18	174 / D1	09/10/J01	FICON EXP16S+ SX	0.10(S)	Availability(S)

Figure 4-46 CHPID to Port Report, sorted by location

4.11.3 CHPID to CU Report

This report is created in way that is similar to the CHPID Report. Click **Preview Report** → **CHPID to Control Unit Report**. The CHPID Mapping Tool displays the CHPID to Control Unit (CU) Report in a Report tab within the CMT (see Figure 4-47).

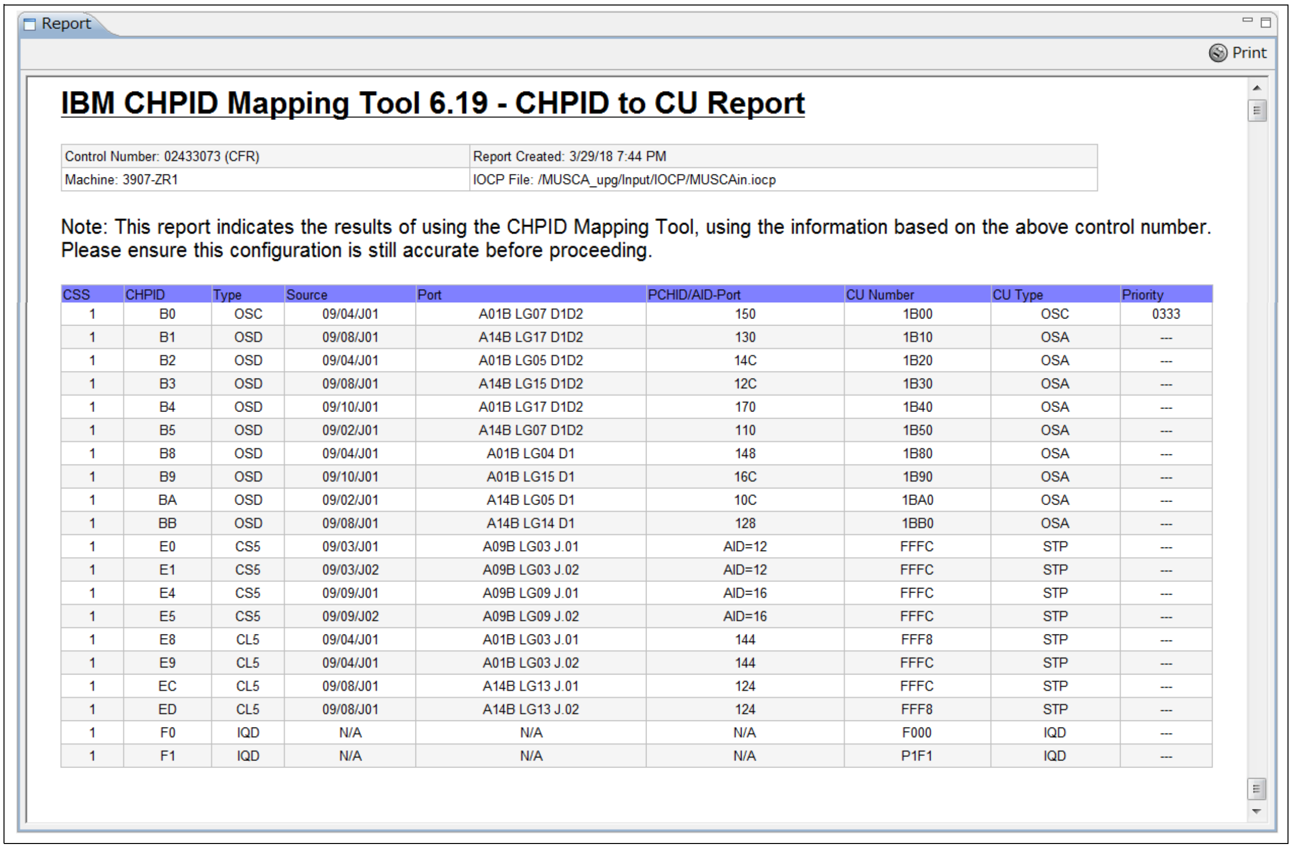

IBM CHPID Mapping Tool 6.19 - CHPID to CU Report

Control Number: 02433073 (CFR)	Report Created: 3/29/18 7:44 PM
Machine: 3907-ZR1	IOCP File: /MUSCA_upg/Input/IOCP/MUSCAin.iocp

Note: This report indicates the results of using the CHPID Mapping Tool, using the information based on the above control number. Please ensure this configuration is still accurate before proceeding.

CSS	CHPID	Type	Source	Port	PCHID/AID-Port	CU Number	CU Type	Priority
1	B0	OSC	09/04/J01	A01B LG07 D1D2	150	1B00	OSC	0333
1	B1	OSD	09/08/J01	A14B LG17 D1D2	130	1B10	OSA	---
1	B2	OSD	09/04/J01	A01B LG05 D1D2	14C	1B20	OSA	---
1	B3	OSD	09/08/J01	A14B LG15 D1D2	12C	1B30	OSA	---
1	B4	OSD	09/10/J01	A01B LG17 D1D2	170	1B40	OSA	---
1	B5	OSD	09/02/J01	A14B LG07 D1D2	110	1B50	OSA	---
1	B8	OSD	09/04/J01	A01B LG04 D1	148	1B80	OSA	---
1	B9	OSD	09/10/J01	A01B LG15 D1	16C	1B90	OSA	---
1	BA	OSD	09/02/J01	A14B LG05 D1	10C	1BA0	OSA	---
1	BB	OSD	09/08/J01	A14B LG14 D1	128	1BB0	OSA	---
1	E0	CS5	09/03/J01	A09B LG03 J.01	AID=12	FFFC	STP	---
1	E1	CS5	09/03/J02	A09B LG03 J.02	AID=12	FFFC	STP	---
1	E4	CS5	09/09/J01	A09B LG09 J.01	AID=16	FFFC	STP	---
1	E5	CS5	09/09/J02	A09B LG09 J.02	AID=16	FFFC	STP	---
1	E8	CL5	09/04/J01	A01B LG03 J.01	144	FFF8	STP	---
1	E9	CL5	09/04/J01	A01B LG03 J.02	144	FFFC	STP	---
1	EC	CL5	09/08/J01	A14B LG13 J.01	124	FFFC	STP	---
1	ED	CL5	09/08/J01	A14B LG13 J.02	124	FFF8	STP	---
1	F0	IQD	N/A	N/A	N/A	F000	IQD	---
1	F1	IQD	N/A	N/A	N/A	P1F1	IQD	---

Figure 4-47 CHPID to CU Report

4.12 Creating an updated IOCP

Now we need to create a CMT updated IOCP statements file that must be imported back into the IODF by using HCD. This IOCP statements file now includes CHIDs that are assigned to CHPIDs.

To create the IOCP, complete the following steps:

1. Select **File** → **Export IOCP input file** (see Figure 4-48).

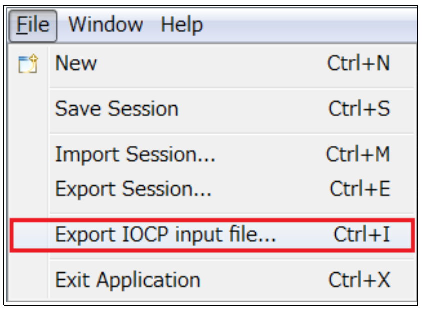

Figure 4-48 Export IOCP input file

2. Enter the Export Path and IOCP Name for the IOCP output file and click **Finish** (see Figure 4-49).

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

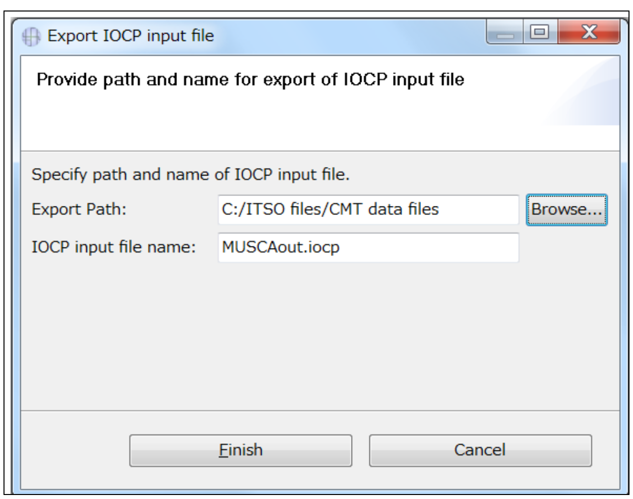

Figure 4-49 Export IOCP File

3. Select **File** → **Save Session** (see Figure 4-50).

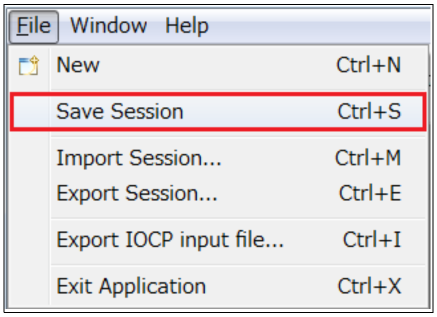

Figure 4-50 Save session

Note: Consider saving your project before exiting the CMT application.

4.13 HCD: Updating the 3907 work IODF with CHIDs

After you map the CHIDs to CHPIDs by using the CHPID Mapping Tool, transfer this information back into HCD. In our example, we call the host file name `‘ITSO1.IODF78.IOCPOUT.MUSCA’`.

To update the IODF with the CHIDs, complete the following steps:

1. Upload the IOCP file that was created by the CMT (`MUSCAout.iocp`, in the example) to the z/OS image. Use a file transfer facility, such as the one in IBM Personal Communications or an equivalent FTP program. Be sure to use TEXT as the transfer type and allocate the z/OS file with RECFM=F or FB and LRECL=80.

   In the updated IOCP statements file, notice that the CMT keeps a reference to the CCN. Also, note the CU Priority values that were added for the OSC control units.

Remember: Control unit priorities are stored in the IOCP output file that is created by CMT. This file is migrated back into HCD. HCD maintains these priorities and outputs them when it creates another IOCP deck. These files are in the form of commented lines at the end of the IOCP deck (see Example 4-2).

Example 4-2 Updated IOCP statements file: With CMT statements

```
        CNTLUNIT CUNUMBR=FFFC,PATH=((CSS(2),E1,E0,E3,E2,E5,E4)),      *
              UNIT=CFP
         IODEVICE ADDRESS=(FFA5,007),CUNUMBR=(FFFC),UNIT=CFP
         IODEVICE ADDRESS=(FFBD,007),CUNUMBR=(FFFC),UNIT=CFP
         IODEVICE ADDRESS=(FFC4,008),CUNUMBR=(FFFC),UNIT=CFP
         IODEVICE ADDRESS=(FFCC,008),CUNUMBR=(FFFC),UNIT=CFP
         IODEVICE ADDRESS=(FFD4,008),CUNUMBR=(FFFC),UNIT=CFP
         IODEVICE ADDRESS=(FFDC,008),CUNUMBR=(FFFC),UNIT=CFP
*CMT* VERSION=000
*CMT* CCN=02433073(CFR from ResourceLink)
*CMT* 1B00.0=0333,1B00.1=0333,PF100.0=0001,PF200.0=0001,PF101.0=0001
*CMT* PF201.0=0001,PF102.0=0001,PF202.0=0001,PF103.0=0001,PF203.0=0001
*CMT* PFD0.0=0001,PFC0.0=0001,PFD1.0=0001,PFB0.0=0001,PFC1.0=0001
*CMT* PFA0.0=0001,PFB1.0=0001,PFA1.0=0001,PFB2.0=0001,PFA2.0=0001
*CMT* PFB3.0=0001,PFA3.0=0001
******************************** Bottom of Data ********************************
```

2. From the HCD main panel, enter the work IODF name used. Select option **5. Migrate configuration data** (see Figure 4-51).

```
                       z/OS V2.3 HCD
Command ===> ________________________________________________________________

                           Hardware Configuration

Select one of the following.

5   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 . . . 'ITSO1.IODF78.WORK'                    +
```

Figure 4-51 Hardware Configuration: Migrate configuration data

3. From the Migrate Configuration Data panel (see Figure 4-52), select option **1. Migrate IOCP/OS data** and press Enter.

```
+----------- Migrate Configuration Data ------------+
|                                                   |
|                                                   |
| Select one of the following tasks.                |
|                                                   |
| 1_  1.  Migrate IOCP/OS data                      |
|     2.  *Migrate switch configuration data        |
|                                                   |
| * = requires TSA I/O Operations                   |
|                                                   |
|  F1=Help     F2=Split    F3=Exit     F9=Swap      |
| F12=Cancel                                        |
+---------------------------------------------------+
```

Figure 4-52 Migrate Configuration Data: Migrate IOCP/OS data

4. The Migrate IOCP Data panel opens (see Figure 4-53). Complete the following fields and then, press Enter:

Processor ID Use the same ID used to create the IOCP input deck.

OS configuration ID This configuration is the operating system configuration that is associated with the processor.

IOCP only input data set This data set is specified when the `MUSCAout.iocp` file was uploaded to z/OS.

Processing mode Select option 2 to save the results of the migration. However, before using option 2, try to migrate by using option 1 to validate the operation.

Migrate options Select option 3 for PCHIDS. Only the PCHIDs are migrated into the work IODF.

```
+-------------------- Migrate IOCP / MVSCP / HCPRIO Data ---------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Processor ID . . . . . . . . . . . . MUSCA___  +   CSS ID . . . . . . _  +  |
| OS configuration ID  . . . . . . . . ITSOTEST  +                            |
|                                                                             |
| Combined IOCP/MVSCP input data set . _____________________________________  |
| IOCP only input data set . . . . . . 'ITSO1.IODF78.IOCPOUT.MUSCA'_________  |
| MVSCP only or HCPRIO input data set  _____________________________________  |
|            Associated with processor ________  +                            |
|                            partition ________  +                            |
| Processing mode  . . . . . . . . . . 2   1.  Validate                       |
|                                          2.  Save                           |
|                                                                             |
| Migrate options  . . . . . . . . . . 3   1.  Complete                       |
|                                          2.  Incremental                    |
|                                          3.  PCHIDs                         |
| MACLIB used  . . . . . . . 'SYS1.MACLIB'                                    |
| Volume serial number . . . ______  + (if not cataloged)                     |
|                                                                             |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F9=Swap    F12=Cancel      |
+-----------------------------------------------------------------------------+
```

Figure 4-53 Migrate IOCP / MVSCP / HCPRIO Data: Data fields to be updated

HCD displays any errors or warning messages that result from the migration action. In the example, the only message generated indicates that the migration was successful (see Figure 4-54).

```
------------------------- Migration Message List ---------------------------
  Query  Help
---------------------------------------------------------------------------
                                                                 Row 1 of 2
Command ===> ___________________________________________ Scroll ===> CSR

Messages are sorted by severity. Select one or more, then press Enter.

/ Statement  Orig Sev Message Text
_                  I   I/O configuration successfully written to the IODF
#                      ITSO1.IODF78.WORK.
***************************** Bottom of data ******************************
```

Figure 4-54 Migration Message List: successful message

The work IODF now contains the CHPID definitions and the mapping to CHIDs that was done by using the CMT.

5. Press PF3. The following message is displayed:

   `IOCP/Operating system deck migration processing complete, return code = 0.`

6. Press PF3 again.

4.14 More steps and processes

For more information about the next steps, see Chapter 5, “Production IODF and setting up the CPC” on page 85.

5

Production IODF and setting up the CPC

This chapter includes the following topics:

5.1 Building the 3907 production IODF

To use the definitions that were updated in HCD, create a 3907 production IODF from your work IODF. Then, remotely or locally write the IODF to the 3907 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** (see Figure 5-1).

```
                       z/OS V2.3 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 . . . 'ITSO1.IODF78.WORK'                   +h
```

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

2. The Activate or Process Configuration Data panel is displayed (see Figure 5-2). Select option **1. Build production I/O definition file** and 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.  *Activate switch configuration              |
|     9.  *Save switch configuration                  |
|     10. Build I/O configuration data                |
|     11. Build and manage System z cluster IOCDSs,   |
|         IPL attributes and dynamic I/O changes      |
|     12. Build validated work I/O definition file    |
|                                                     |
| * = requires TSA I/O Operations                     |
|  F1=Help     F2=Split    F3=Exit     F9=Swap        |
| F12=Cancel                                          |
+-----------------------------------------------------+
```

Figure 5-2 Activate or Process Configuration Data: Build production IODF

3. HCD displays the Message List panel (see Figure 5-3). Verify that only severity “W” (warning) messages are present and that they are normal for the configuration. Correct any other messages that should not occur and attempt to build the production IODF again. Continue this process until no messages occur 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
| #              group(s): 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
| #              group(s): 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
| #              group(s): 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 production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (see Figure 5-4). Enter 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 . . . : 'ITSO1.IODF78.WORK'                         |
|                                                                    |
| Production IODF name . 'SYS6.IODF79'______________________         |
| 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                    |
|                                                                    |
|                                                                    |
|                                                                    |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F9=Swap   F12=Cancel |
+--------------------------------------------------------------------+
```

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  . : 'SYS6.IODF79'                           |
|                                                                   |
| Descriptor field 1  . . . SYS6                                    |
| Descriptor field 2  . . . IODF79                                  |
|  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 SYS6.IODF79 created.
```

Proceed to the next section to implement the configuration on the 2965 in preparation for its upgrade to a 3907.

5.2 Writing the IOCP to the 2965 processor using HCD

Now that a production IODF that is named `SYS6.IODF79` is set up, you can write the IOCP data from the IODF to the IOCDS on the server that you want to upgrade (for example, LEPUS).

The IOCDS is 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 as described in 5.1, “Building the 3907 production IODF” on page 86. Then, press Enter.

```
                        z/OS V2.3 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 . . . 'SYS6.IODF79'                     +
```

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 System z cluster IOCDSs, IPL attributes and dynamic I/O changes**. 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.  *Activate switch configuration              |
|     9.  *Save switch configuration                  |
|     10. Build I/O configuration data                |
|     11. Build and manage System z cluster IOCDSs,   |
|         IPL attributes and dynamic I/O changes      |
|     12. Build validated work I/O definition file    |
|                                                     |
| * = requires TSA I/O Operations                     |
|  F1=Help     F2=Split    F3=Exit     F9=Swap        |
| F12=Cancel                                          |
+-----------------------------------------------------+
```

Figure 5-7 Build and manage System z cluster IOCDSs, IPL attributes, and dynamic I/O changes

In this example, it is assumed that connectivity exists to the 2965 that is being upgraded over the HMC local area network to write an IOCDS.

If the server that is 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 CMT.

> **Tip:** The Support Element can now read an IOCP file that was written to a USB flash memory drive.

3. The System z Cluster List panel opens (see Figure 5-8). In the list, select the 2965 that is being upgraded by entering a forward slash (/) to update one of its IOCDSs. Then, press Enter.

```
                            System z Cluster List                   Row 1 of 3
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
_ IBM390PS.CETUS     3906   M03     CETUS
/ IBM390PS.LEPUS     2965   N20     MUSCB
_ IBM390PS.MUSCA     3907   ZR1     MUSCA
******************************* Bottom of data *******************************
```

Figure 5-8 System z Cluster List: Selecting processor for IOCDS replace

4. The Actions on selected CPCs panel opens (see Figure 5-9). Select option **1. Work with IOCDSs** and 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)             |
|                                                              |
|  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 2965 replacement by entering a forward slash (/) next to it. Then, press Enter.

```
                                  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.MUSCB    IODF78   LPAR   Alternate  No        No          No
_ A1.MUSCB    IODF79   LPAR   POR        Yes       No          Yes-POR
_ A2.MUSCB    IODF72   LPAR   Alternate  No        No          No
/ A3.MUSCB    IODF74   LPAR   Alternate  No        No          No
******************************* Bottom of data *******************************
```

Figure 5-10 IOCDS List: Selecting IOCDS for replace

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)            |
|                                                             |
|                                                             |
|  F1=Help     F2=Split    F3=Exit     F9=Swap    F12=Cancel  |
+-------------------------------------------------------------+
```

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 press Enter.

```
+------------------------------ Build IOCDSs -------------------------------+
|                                                              Row 1 of 1 |
| Command ===> ___________________________________________ Scroll ===> CSR  |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| IODF name . . . . . . . . . : 'SYS6.IODF79'                               |
|                                                                           |
| Title1 . IODF79__________________________________________________________ |
| Title2 : SYS6.IODF79 - 2018-03-15 10:36                                   |
|                                                                           |
|                            Write IOCDS in                                 |
| IOCDS        Switch IOCDS  preparation of upgrade                         |
| A3.MUSCB     No            Yes                                            |
| ***************************** Bottom of data ****************************** |
|                                                                           |
|                                                                           |
|  F1=Help       F2=Split       F3=Exit       F5=Reset       F7=Backward    |
|  F8=Forward    F9=Swap       F12=Cancel     F22=Command                   |
+---------------------------------------------------------------------------+
```

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 3907 from the existing hardware. The `Yes` value enables writing an IOCDS that contains information that the current hardware does not recognize.

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

```
+------------------------------- Build IOCDSs --------------------------------+
|                                                                             |
+---------- Confirm Write IOCDS in preparation of processor upgrade ----------+
|                                                                Row 1 of 1 |
| Command ===> ___________________________________________ Scroll ===> CSR   |
|                                                                             |
| 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.MUSCB                                                                    |
| ***************************** Bottom of data ****************************** |
|                                                                             |
|                                                                             |
|  F1=Help        F2=Split       F3=Exit        F7=Backward    F8=Forward     |
|  F9=Swap       F12=Cancel     F22=Command                                   |
+-----------------------------------------------------------------------------+
```

Figure 5-13 Build IOCDSs: Confirm write IOCDS)

9. The Job Statement Information panel is displayed (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 3907 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
-CBDA674I  IOCP successfully completed for A3.MUSCB.
```

If you return to HCD option 2.11 and view the IOCDS, notice that the SNA Address remains at `IBM390PS.LEPUS` (see Figure 5-15).

```
                            System z Cluster List                   Row 1 of 3
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
_ IBM390PS.CETUS     3906   M03     CETUS
_ IBM390PS.LEPUS     2965   N20     MUSCB
_ IBM390PS.MUSCA     3907   ZR1     MUSCA
******************************* Bottom of data ********************************
```

Figure 5-15 System z Cluster List: Selecting processor for IOCDS verify

When you select **IBM390PS.LEPUS**, 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 3907 processor.

The status switches when this processor is upgraded to a 3907. The 2965 IOCDS status changes to `Alternate` and the 3907 IOCDSs changes to `Invalid`.

Tip: Generally, rewrite the IOCDS that is written in preparation for the upgrade at your earliest convenience. Subsequent 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.MUSCB    IODF78   LPAR   Alternate    No        No          No
_ A1.MUSCB    IODF79   LPAR   POR          Yes       No          Yes-POR
_ A2.MUSCB    IODF72   LPAR   Alternate    No        No          No
_ A3.MUSCB    IODF79   LPAR   Invalid      No        Yes         No
******************************* Bottom of data ********************************
```

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

5.3 Creating a Reset Profile on the 3907 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 must occur before you create a Reset Profile:

- The 2965 processor was upgraded to a 3907-ZR1.
- A new HMC was installed with the correct driver level to support the 3907, and it is connected to the customer HMC network.
- The upgraded 3907 (in this example, MUSCA) was defined to the new HMC.
- The 2965, now upgraded to a 3907, was Power on Reset with the Diagnostic (DEFAULT) IOCDS.
- A new IOCP was written to the 3907’s IOCDS from the IODF (IODF79) by using HCD Option 2.11.

The 3907 is now ready to be customized with specific customer definitions.

5.3.2 Building the Reset Profile and pointing to required IOCDS

Now that the IOCP file is written to an IOCDS, build a Reset (power-on reset) Profile to point to that IOCDS. This Reset Profile is used to power on reset the new 3907 after it is upgraded and handed over from the IBM System Service Representative.

To build the profile, complete the following steps:

1. Log on by using SYSPROG authority to the HMC workstation that is supplied with the 3907, or use a remote web browser and select the new 3907.
2. Under Systems Management, click **Systems Management** to expand the list.
3. Under Systems Management, click the radio button that is next to the system to select it (in this example, MUSCA).
4. On the Tasks window, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (see Figure 5-17 on page 97).

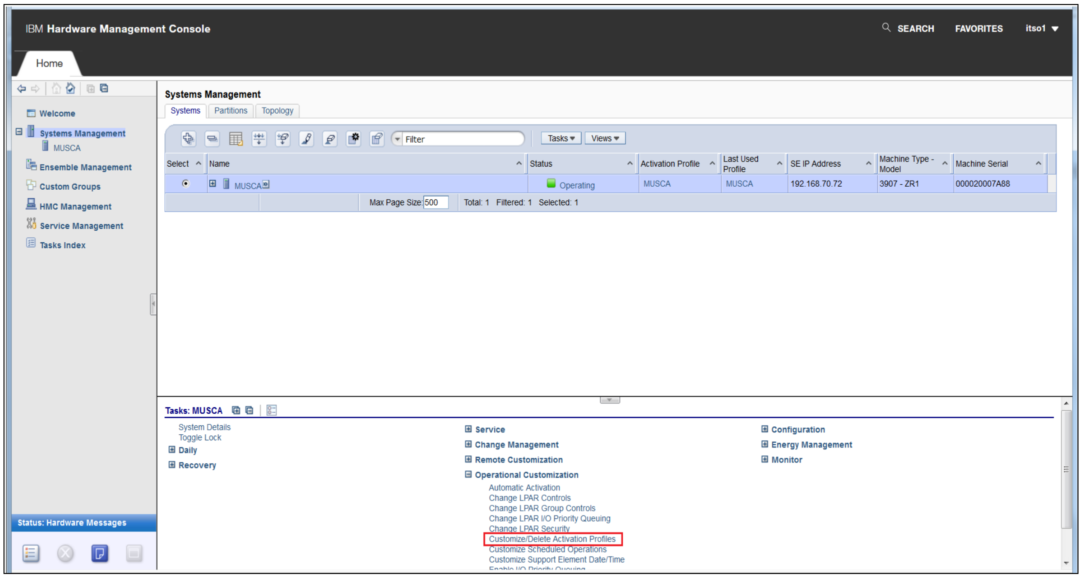

Figure 5-17 Customize Activation Profiles

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 power-on reset 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 is displayed (see Figure 5-18).

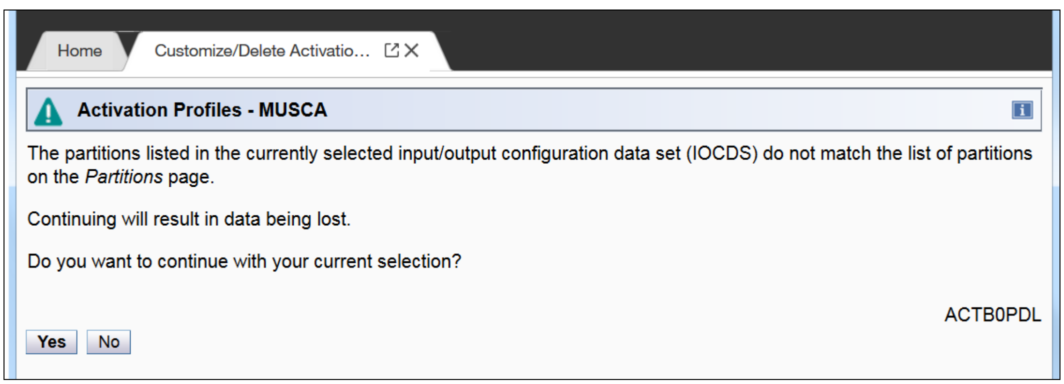

Figure 5-18 Activation Profiles: ACTB0PDL message

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 LPAR names that are defined in the IOCDS that was selected. It also allows you to create Image profiles for those image profiles that cannot be retrieved.
11. 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 on page 98).

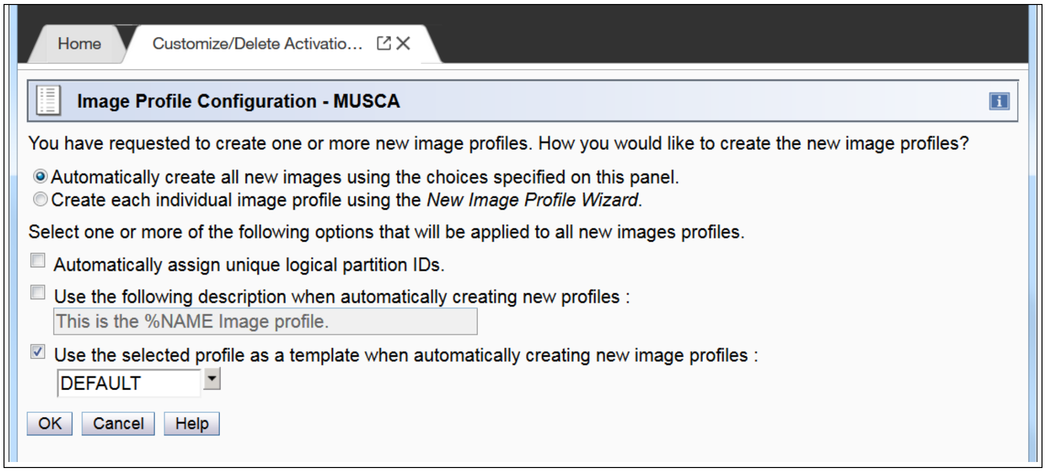

Figure 5-19 Image Profile automatic build options

12. 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).

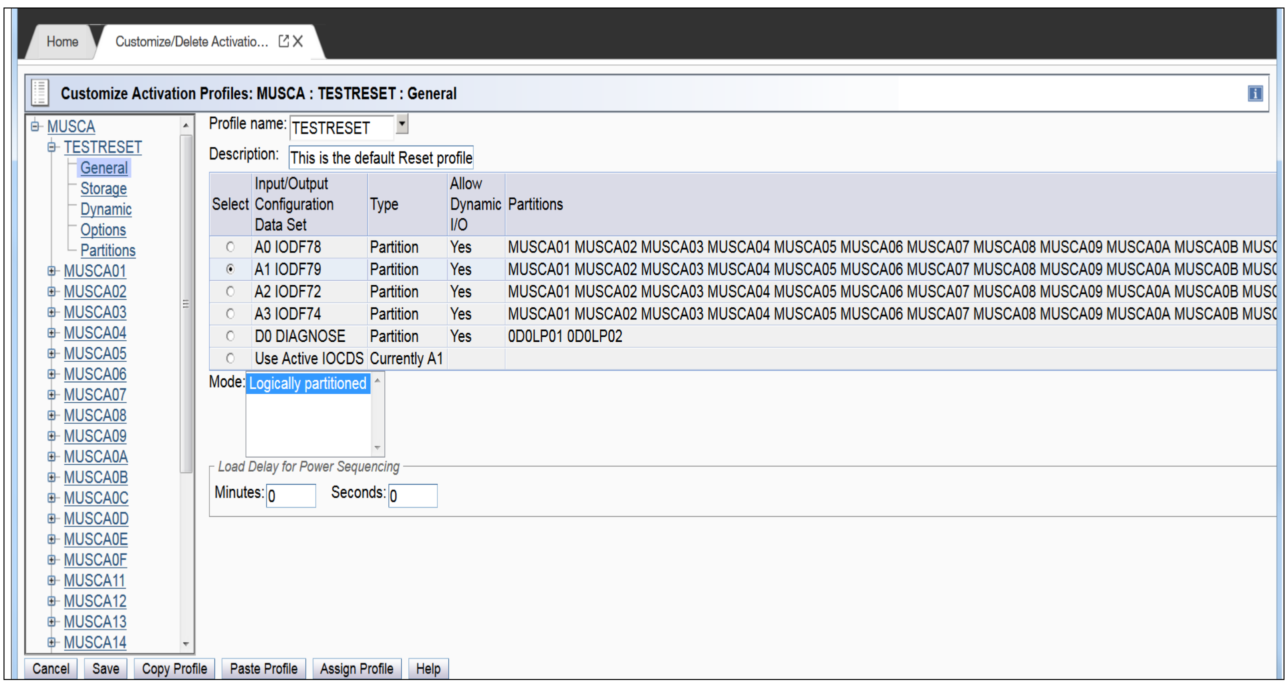

Figure 5-20 Reset and Image profile list: Selecting IOCDS

For information about planning, see Chapter 2, “Planning considerations for CPC in PR/SM mode” on page 7.

5.3.3 Setting up and verifying the Reset Profile

To set up and verify the Reset Profile, complete the following steps:

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

   This display lists all of the partitions that were retrieved by the automatic build for Reset Profile TESTRESET. The partition list also determines all the Image profiles that are activated if the CPC was Power on Reset (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 changed.

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

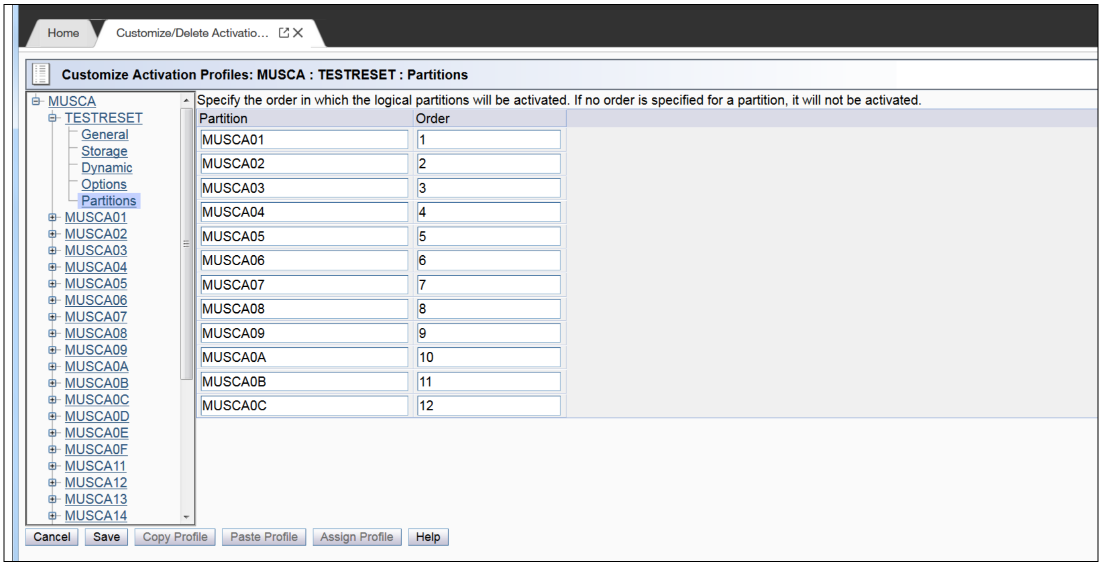

Figure 5-21 Reset and Image profile list: Updating the partition list

3. If any Coupling Facility partitions are defined, HMC asks if 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).

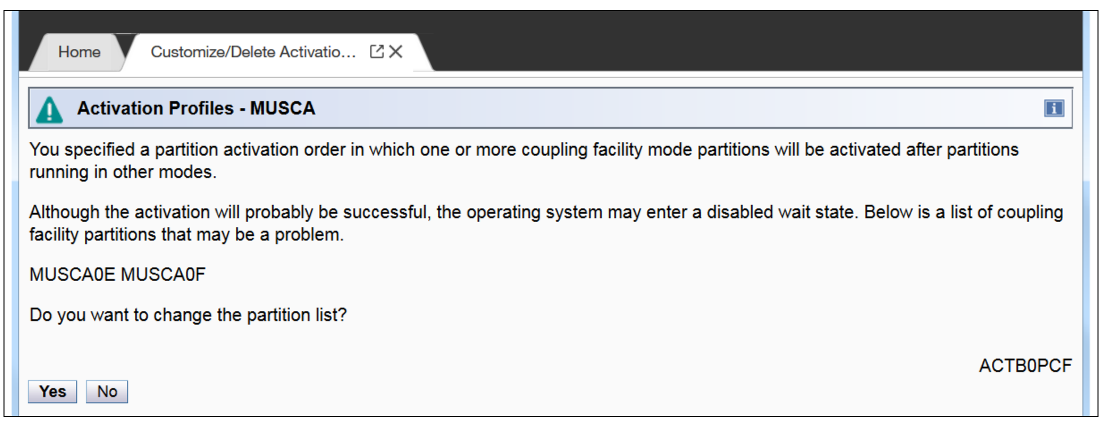

Figure 5-22 Reset and Image profile list: Coupling Facility LPAR verification

5.4 Creating an Image Profile on the 3907 Support Element

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

- General
- Processor
- Security

- Storage
- Options
- Load
- Crypto
- Time Offset

Click one of the Image profiles to set up the partitions parameters. In our example, we select MUSCA01.

5.4.1 Image Profile: General page

The General page is displayed first (see Figure 5-23). Review the following settings:

- Description
- Partition identifier
- Mode
- Clock Type Assignment:
  - Standard time of day
  - Logical partition time offset

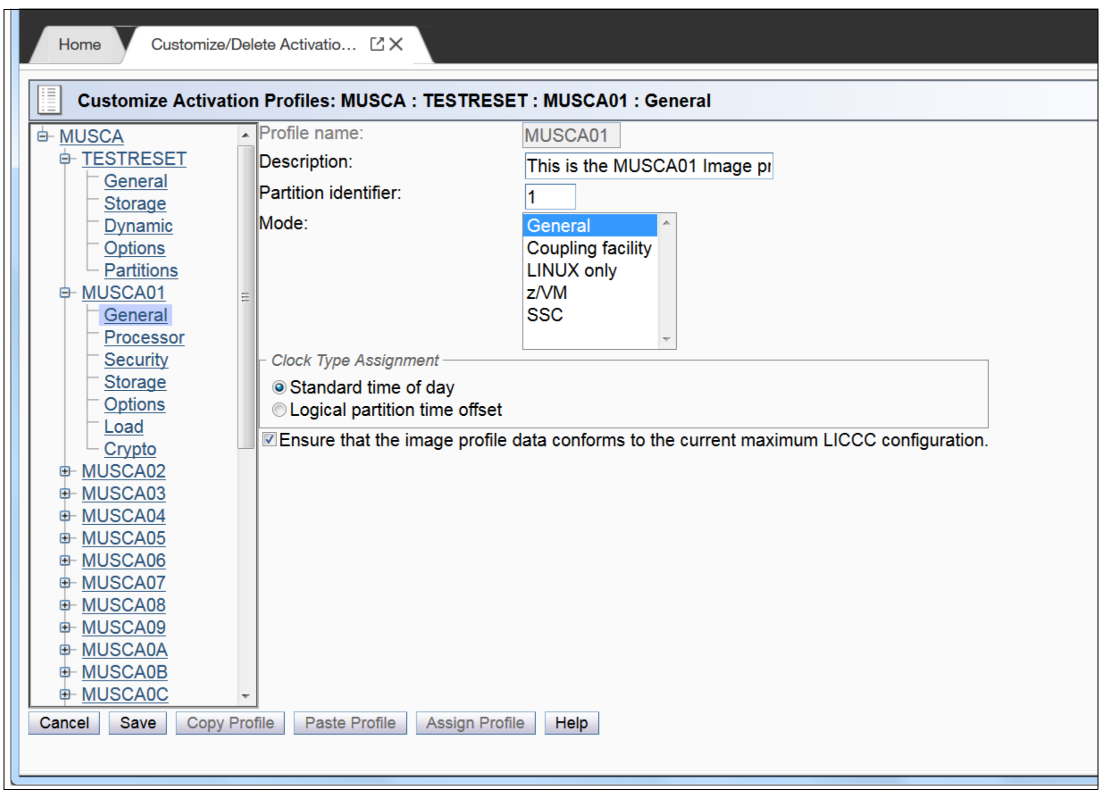

Figure 5-23 Image Profile: General

5.4.2 Image Profile: Processor page

Click the **Processor** link to set up the partitions CPU and weight information (see Figure 5-24). Review the following settings:

- Dedicated processors tick box: Click this setting first if you want to set dedicated CPs, zIIPs, IFLs, or ICFs.
- Central processors (CPs) for Initial and Reserved.
- IBM Z Integrated Information Processors (zIIP) 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.

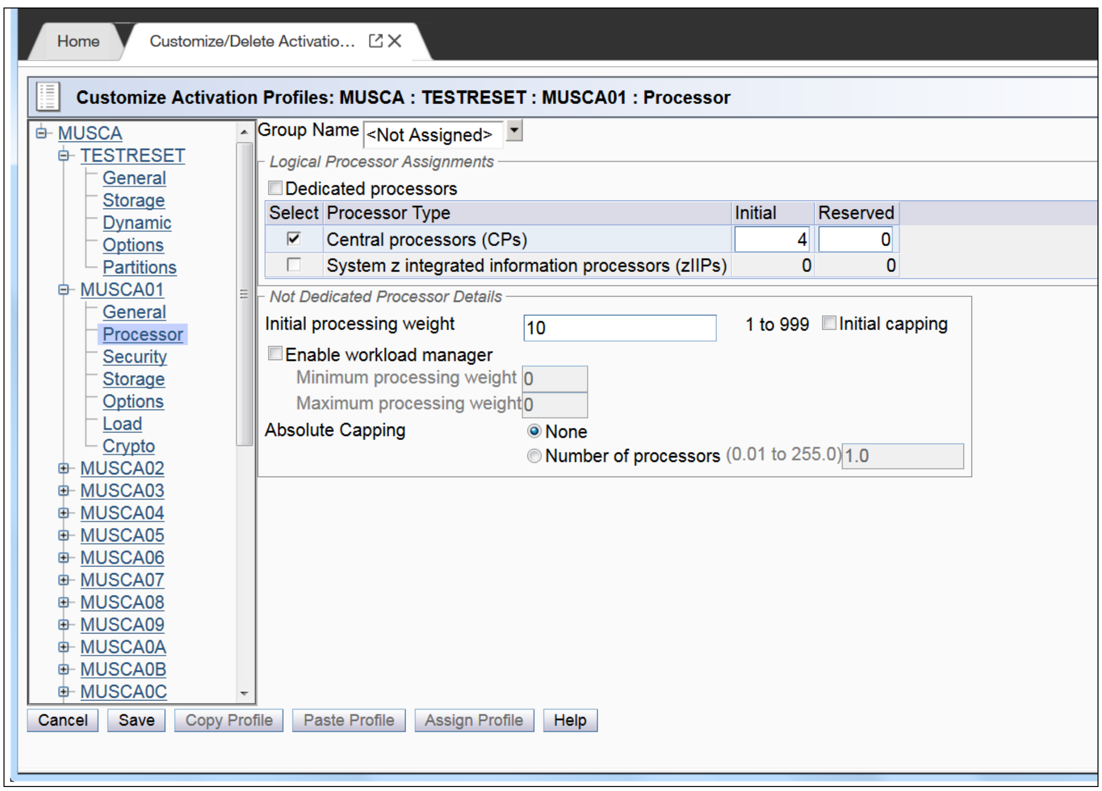

Figure 5-24 Image profile: Processor

5.4.3 Image Profile: Security page

Next, click the **Security** link to set up the partitions 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

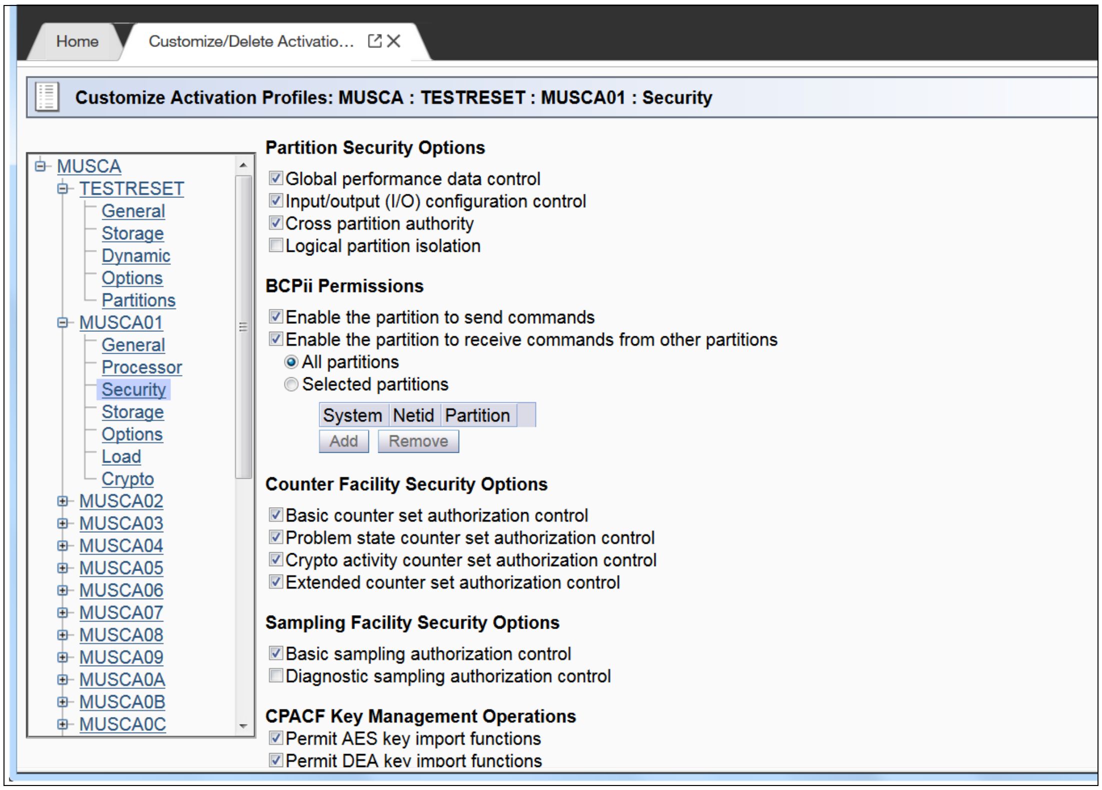

Figure 5-25 Image profile: Security

5.4.4 Image Profile: Storage page

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

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

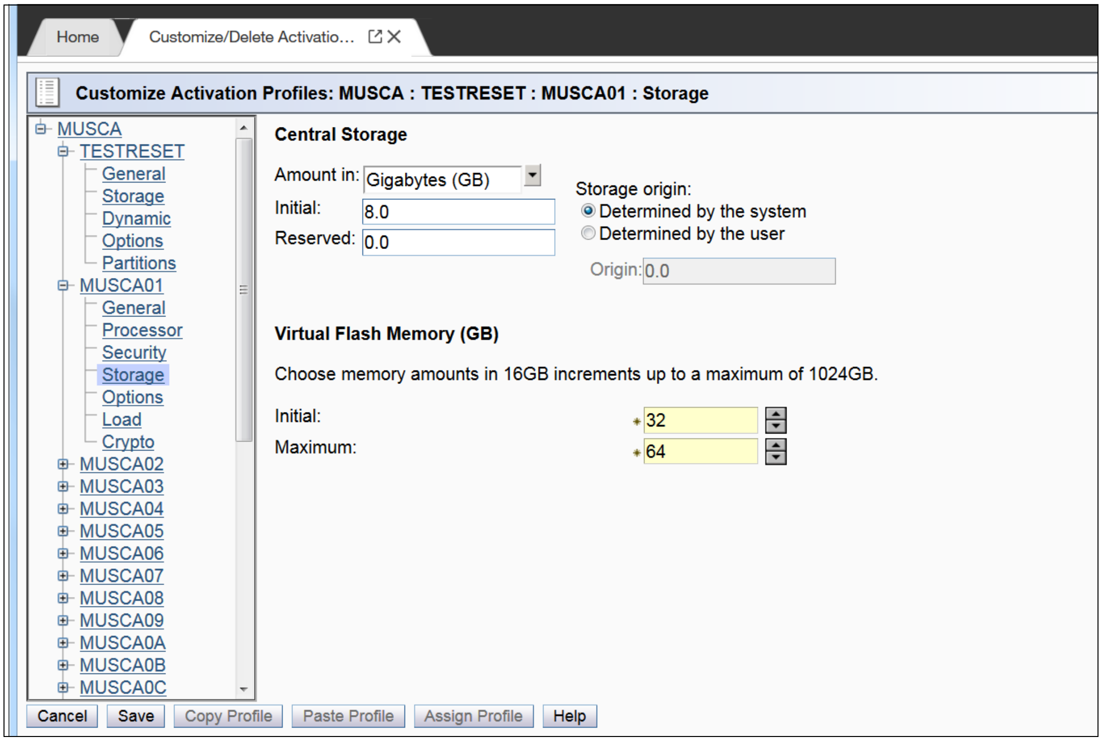

Figure 5-26 Image profile: Storage

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

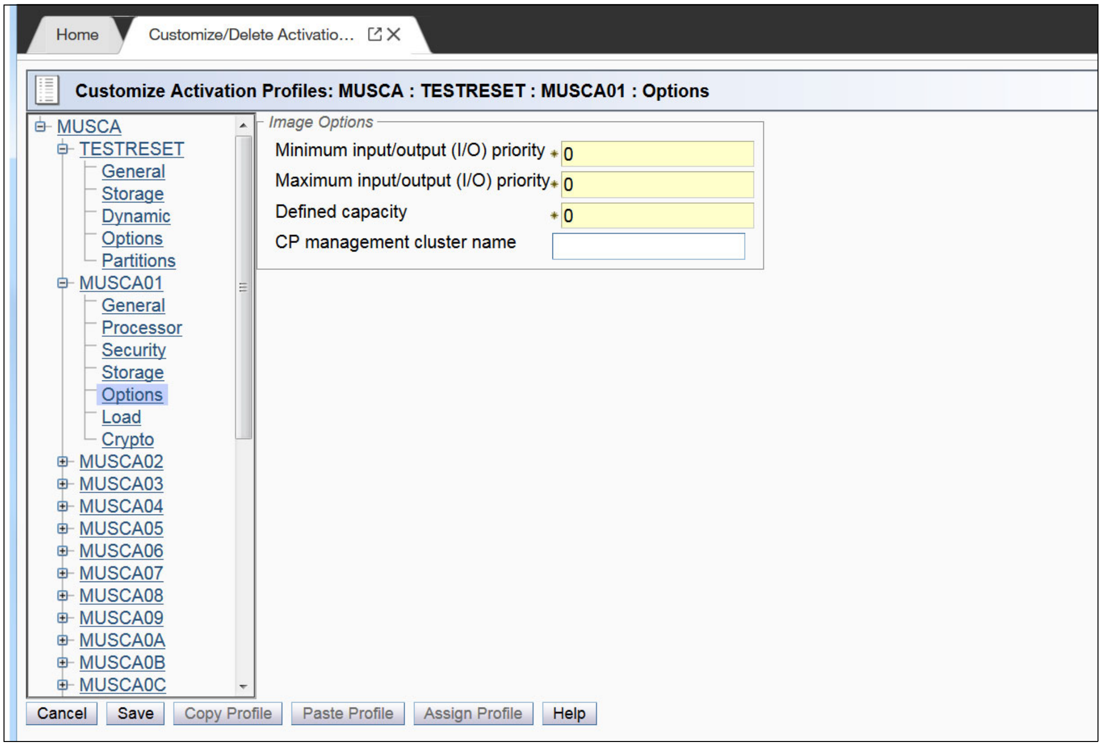

Figure 5-27 Image profile: Options

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 Power-on Reset or Image profile activation (see Figure 5-28). If you prefer not to IPL a z/OS system into a partition during a Power-on Reset or Image profile activation, Load Profiles can be set up and activated and used when required. This concern is described in the following section. Review the following settings:

- Load during activation check box, which determines whether the options in this window are used
- Load type
- Load address
- Time-out value

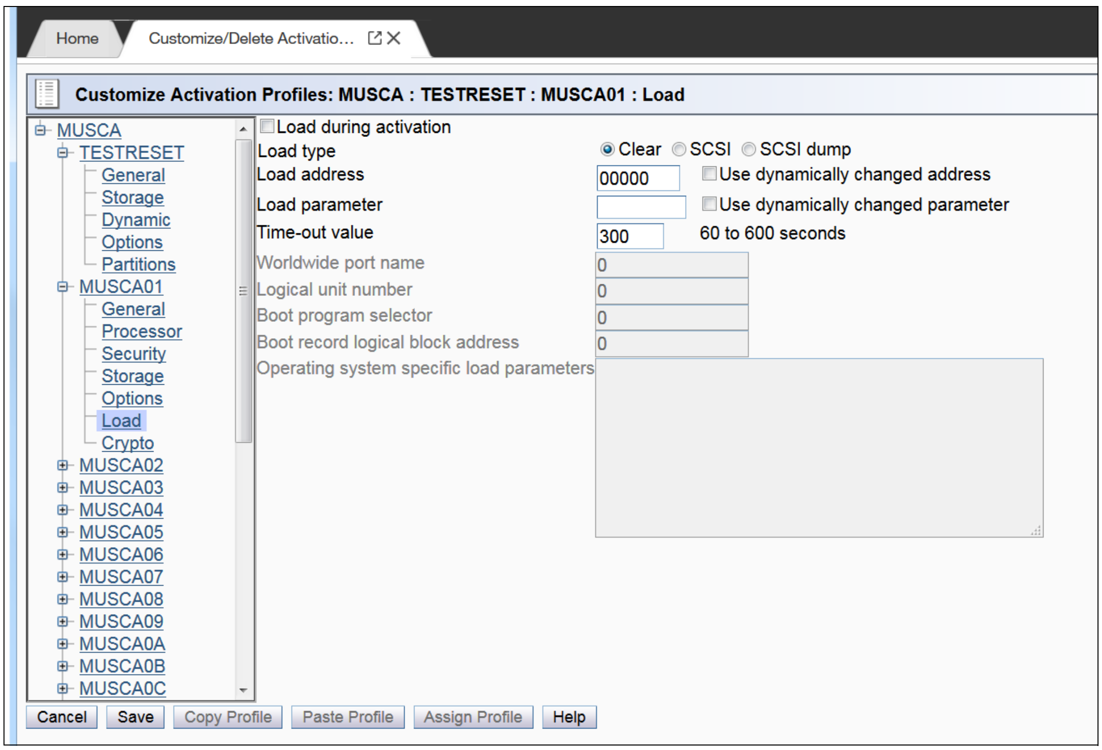

Figure 5-28 Image profile: Load

5.4.7 Image Profile: Crypto page

Click the **Crypto** link where you define the Crypto Domain Index IDs and the number of Crypto engines that are assigned to that Domain ID and whether they are only candidate or also 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.

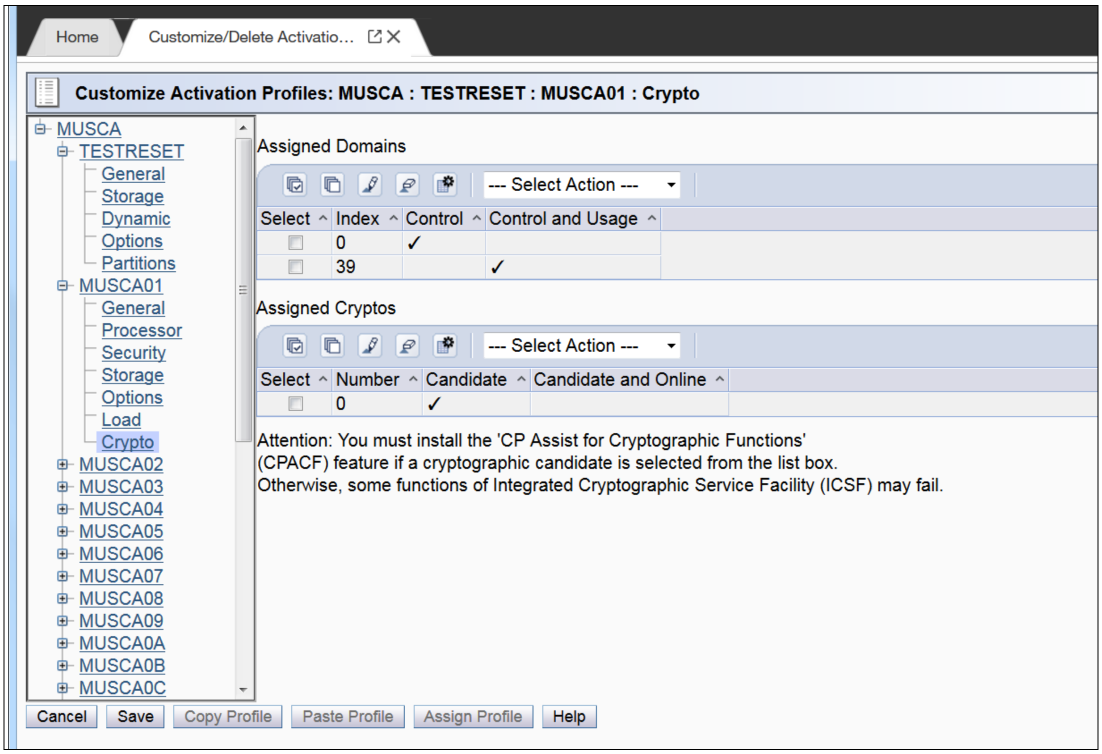

Figure 5-29 Image profile: Crypto

5.4.8 Image Profile: Time Offset

If you selected **Logical partition time offset** in the General window, an extra window is displayed in the Image profile that is named Time Offset. Here, you can select the partitions Time offset against the CPCs time as set by STP (see Figure 5-30). Review the following settings:

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

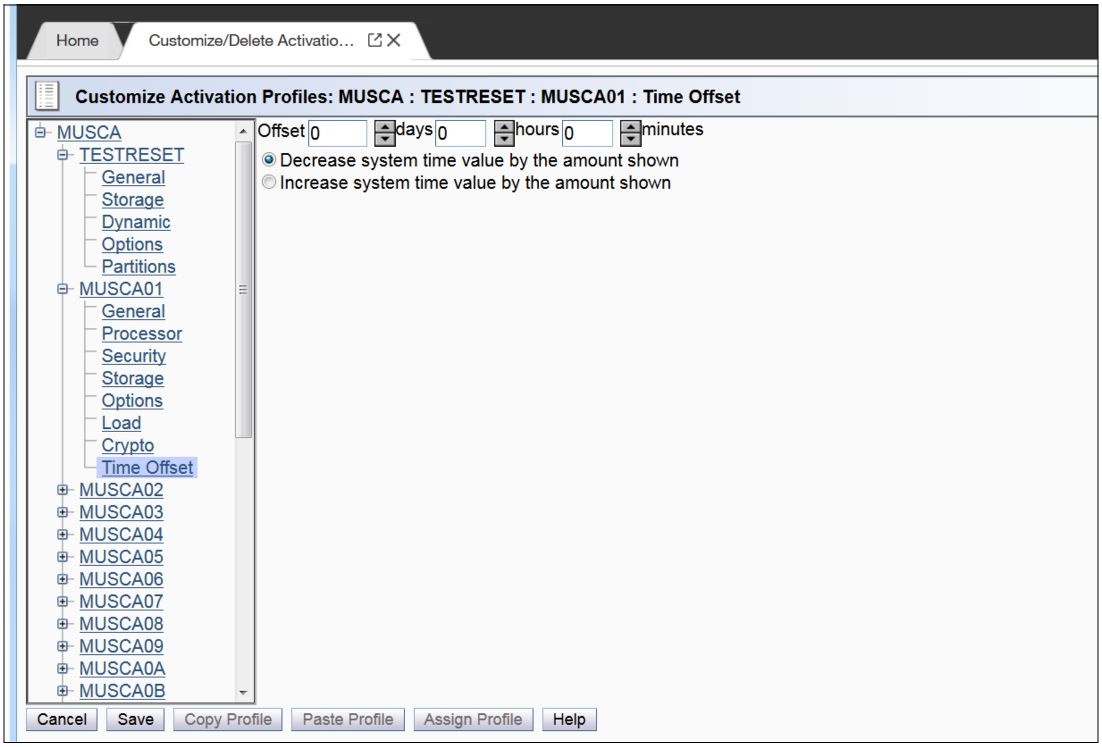

Figure 5-30 Image profile: Time offset

5.4.9 Image Profile: Saving

After you customized all of the required Image profiles for this Reset profile, click **Save** to save the Reset (and Image) profiles for Power-on Reset. If only a few LPARs require activation, it might be easier to deactivate and activate those Image profiles individually. It depends on your situation. HMC prompts you for a confirmation to continue to save. Click **OK** (see Figure 5-31).

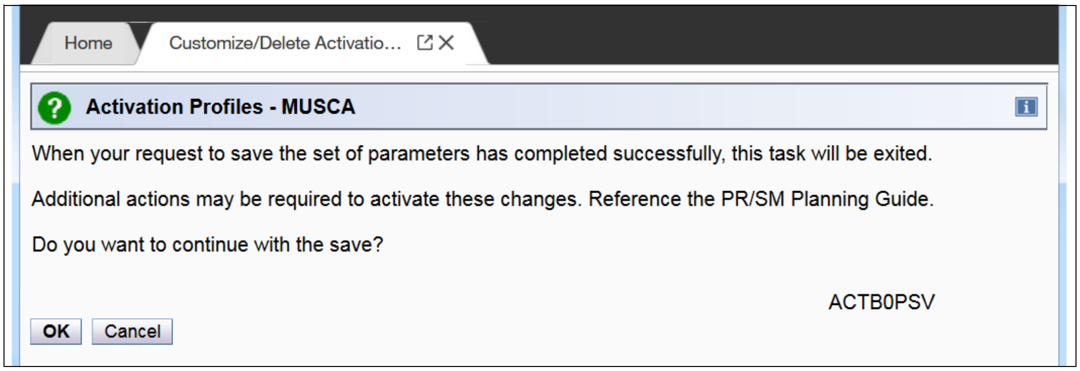

Figure 5-31 Reset and Image profile: Confirmation to save

For more information about planning, see Chapter 2, "Planning considerations for CPC in PR/SM mode" on page 7.

5.5 Performing a Power-On Reset of the 3907

When the 2965 processor is upgraded to a 3907, the IBM System Service Representative performs a POR with a Diagnostic IOCDS.

After this process is complete and the IBM System Service Representative is satisfied with the status of the processor, the IBM System Service Representative hands over the processor to you. You then run another POR by using the Reset Profile that was created as described in 5.3.2, "Building the Reset Profile and pointing to required IOCDS" on page 96.

The 3907 is now ready to be activated (power-on reset) by using the Production Reset Profile. This process is optional but preferred depending on how many partitions that were defined on the processor.

5.5.1 Coupling Facility Links

After the Power-on Reset completes with your specific customer configuration and the coupling links are 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.

This verification is done by displaying the CHPID by using Channel Problem Determination on the HMC.

To access the Channel Problem Determination windows, complete the following steps:

1. Log on by using SYSPROG authority to the HMC for the new 3907.
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, MUSCA).
4. On the Tasks window, click **Recovery** to expand it, and select **Single Object Operations** (see Figure 5-32).

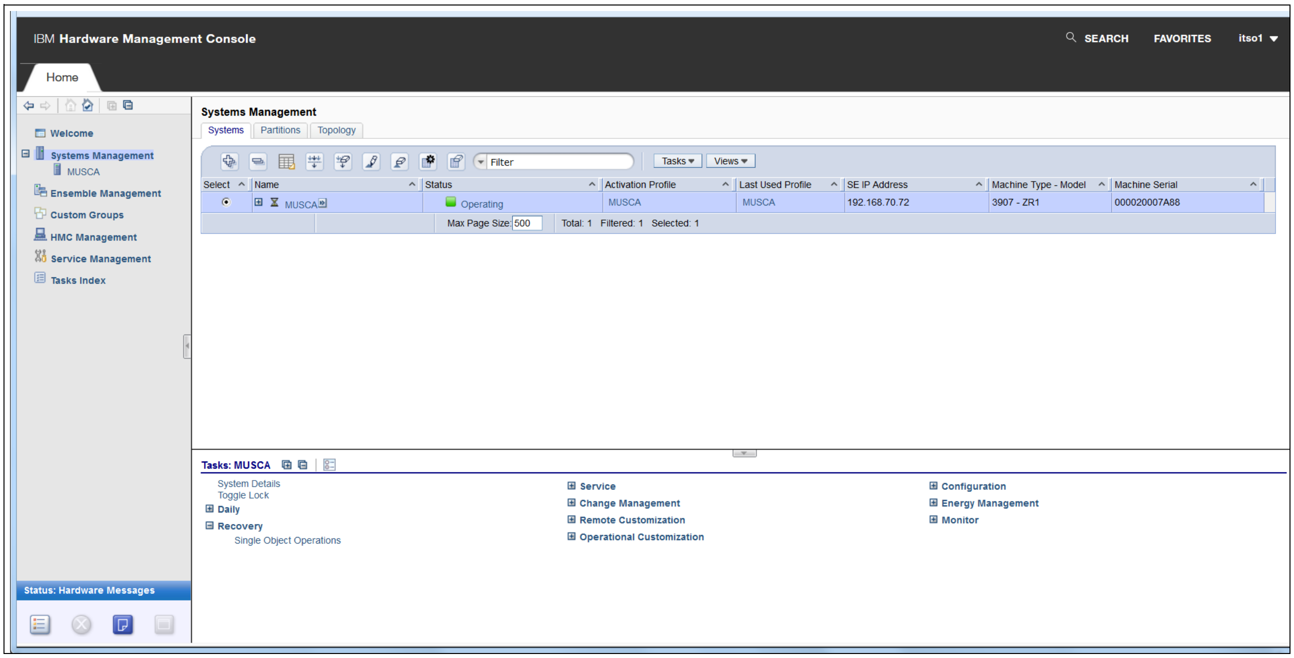

Figure 5-32 Systems Management: Main display

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, MUSCA).
8. Click **Partitions** to expand the list of partitions.
9. Scroll through the list of partitions until you find one of the coupling facility partitions or z/OS partitions to which the coupling links are connected (in our example, we select MUSCA11).
10. Click the partition name to expand the options under the partition name.

11.Click **CHPIDs** to display the CHPID list specific to this LPAR (see Figure 5-33).

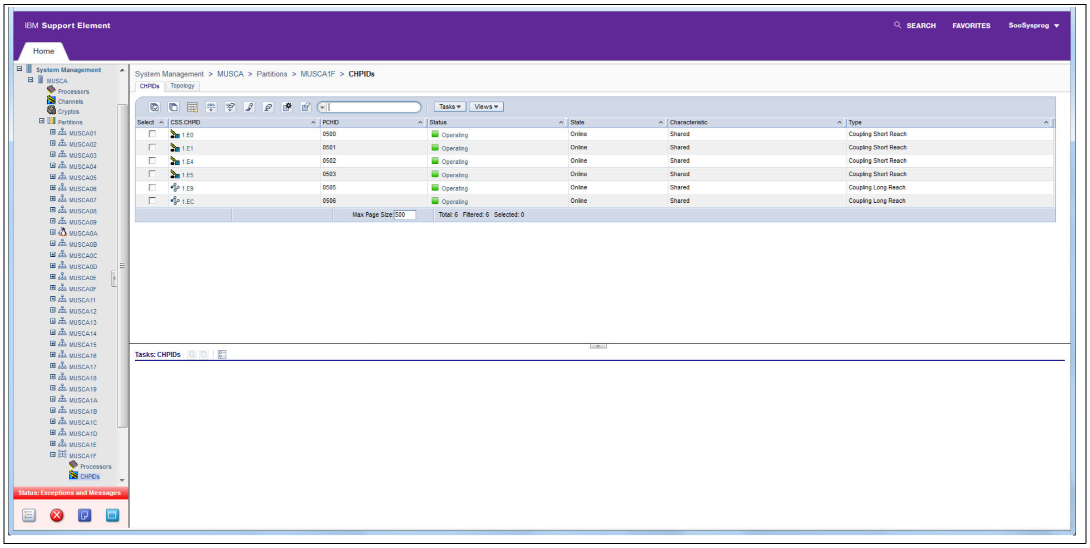

Figure 5-33 System Management: Single Object Operation

12.Select the CHPID that you want to verify to highlight it. We view CHPID **1.EC** (CSS=1, CHPID=EC).

13.The options for this CHPID can be shown by using one of the following methods:

– Click the **>>** symbol that is next to the CHPID to expand out the options.

– Click **CHPID Operations** to expand the options and click **Channel Problem Determination** (see Figure 5-34).

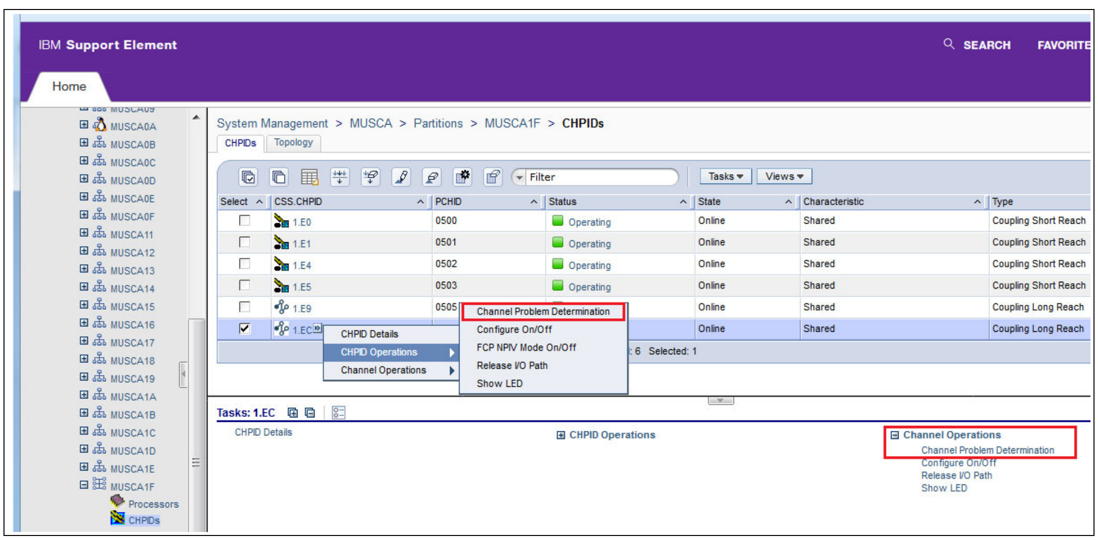

Figure 5-34 System Management: CHPID Operations

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

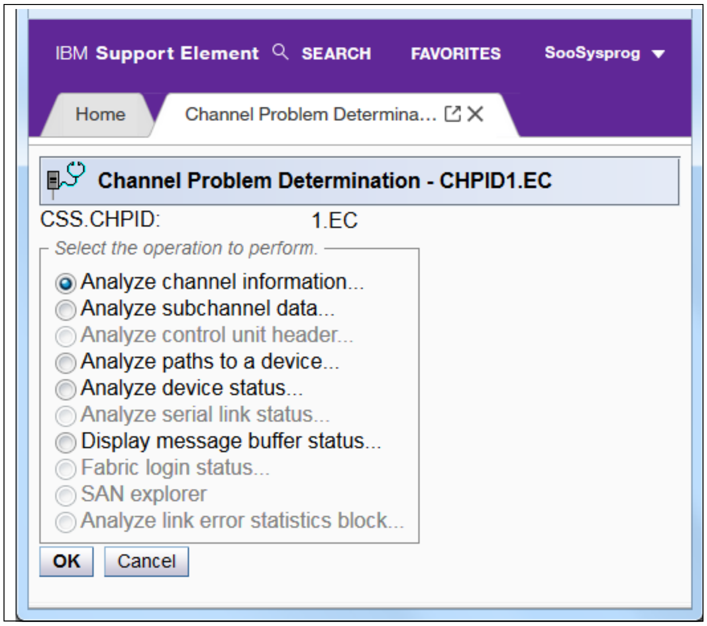

Figure 5-35 Channel Problem Determination: Analyze channel information

15. Observe the following information (see Figure 5-36 on page 112):
   - State: Online
   - Status: Operating
   - Node type: Attached
   - Node status: Valid
   - Type/model: 3906-M03 (device to which the CHPID is connected)
   - Seq. number: EE0F7 (serial number of the device to which the CHPID is connected)
   - Tag: ED (in this case, the destination CHPID of CHPID EC)

   Observe also the PCHID of 0506. This PCHID number is allocated by the CPC when this specific CHPID (Coupling Express LR) is defined to the HSA configuration.

T

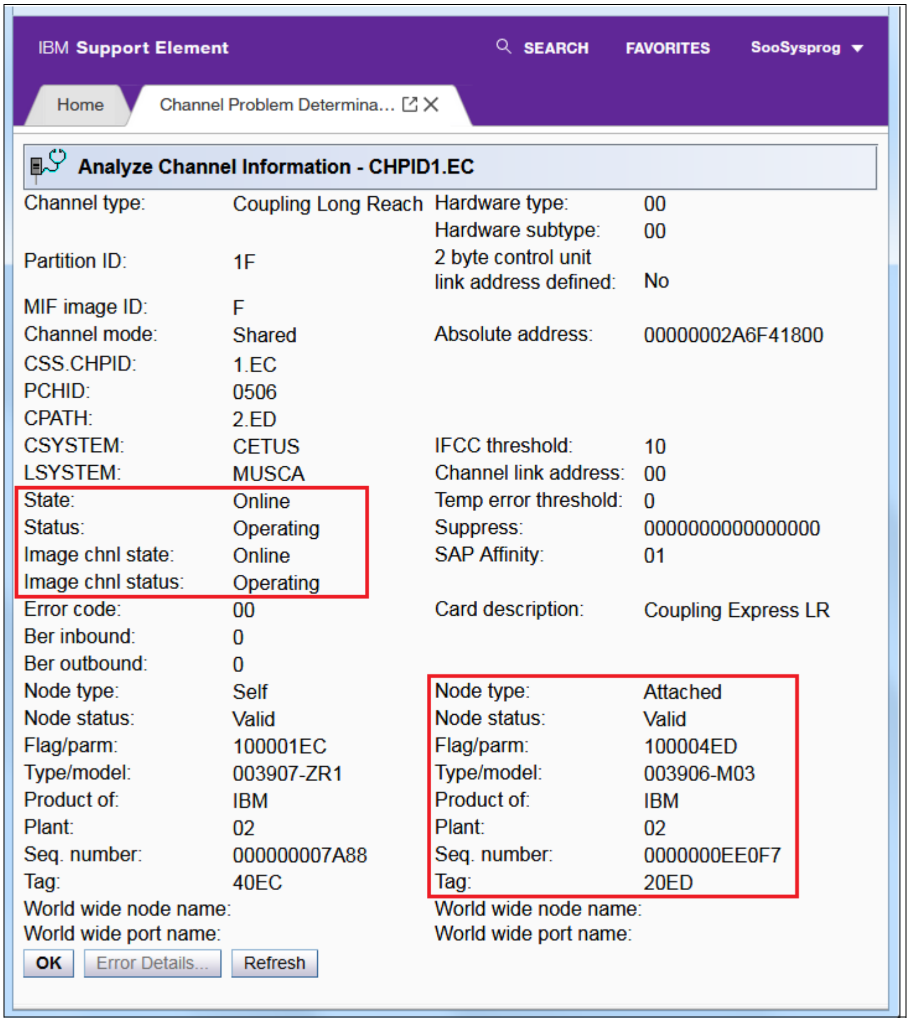

Figure 5-36 Channel Problem Determination: display

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

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

5.5.2 Server Time Protocol configuration

Now that the Coupling Facility links are verified as connected and online, you can set up the Server Time Protocol (STP) configuration.

With z14 ZR1, the STP or Manage System Time option on the HMC under Configuration changed to a graphic user interface (GUI) format.

For more information about the new GUI and how to set up the STP Coordinated Time Network (CTN), see Chapter 8, “Preparing for Sysplex and configuring Server Time Protocol” on page 159.

This section accesses the System (Sysplex) Time option under Single Object Operations for CPC MUSCA to display and verify some information.

To access the System (Sysplex) Time windows under Single Object Operations, complete the following steps:

1. Log on by using SYSPROG authority to the HMC for the new 3907.
2. Click **Systems Management** to expand the list.
3. Under Systems Management, click the radio button that is next to the system to select it (in this example, MUSCA).
4. On the Tasks window, click **Recovery** to expand it, and select **Single Object Operations** (see Figure 5-37).

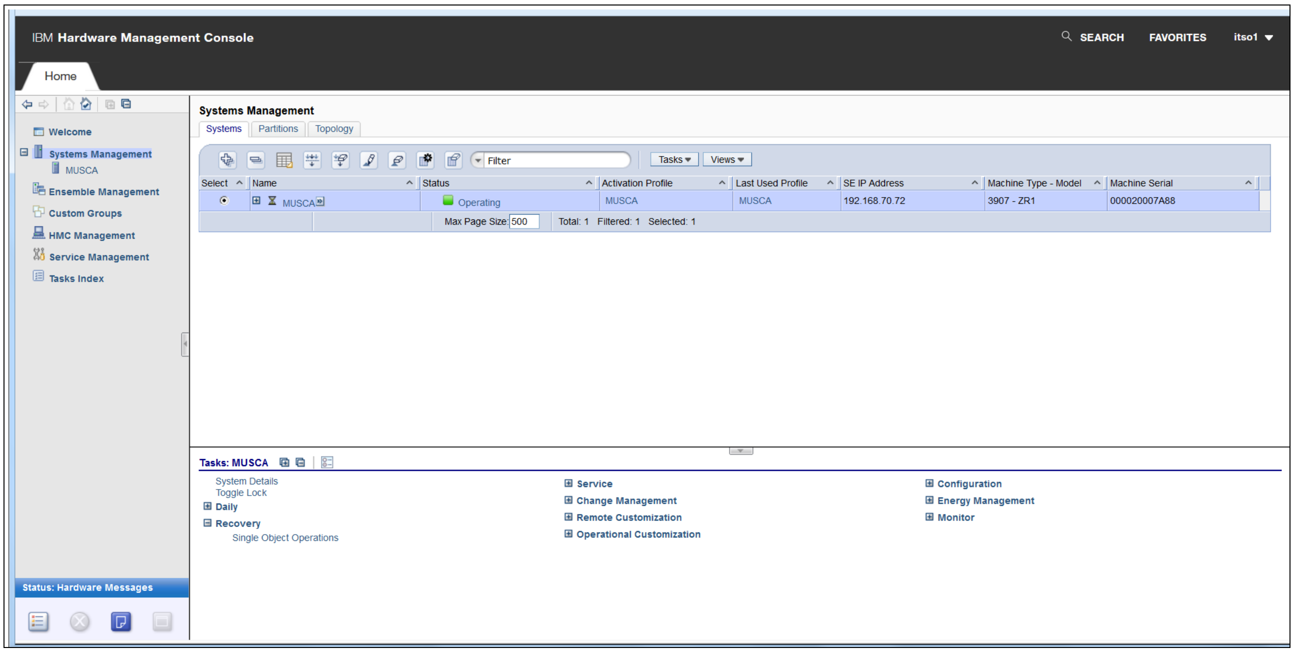

Figure 5-37 Systems Management: Main display

5. Click **System Management** to expand the list.
6. Under Systems Management, click the CPC name to expand the options (in this example, MUSCA).
7. Click **Configuration** to expand the items under the configuration category.

8. Click **System (Sysplex) Time** (see Figure 5-38).

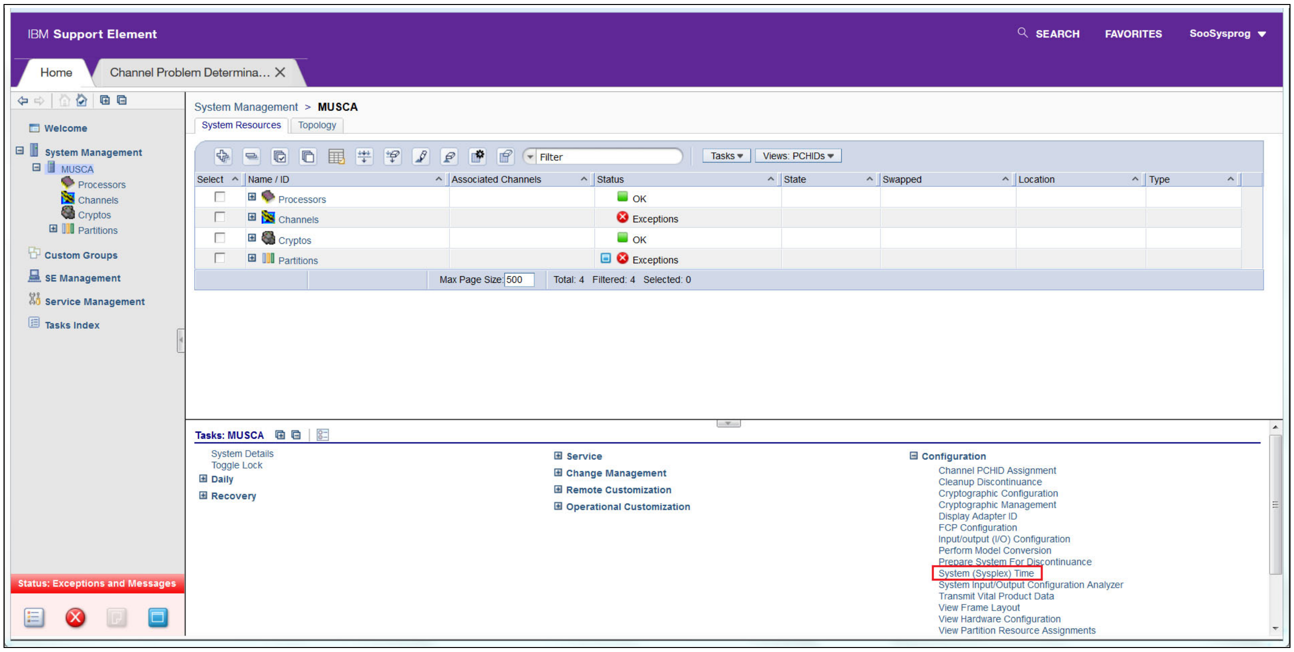

Figure 5-38 Single Object Operations: System (Sysplex) Time

9. Click **Yes** for the Attention window that warns you that any action in the following windows might affect the current LPARs that are Operating.

10. Click the **ETS Configuration** tab to check whether any External Time Source (ETS) definitions exist (see Figure 5-39).

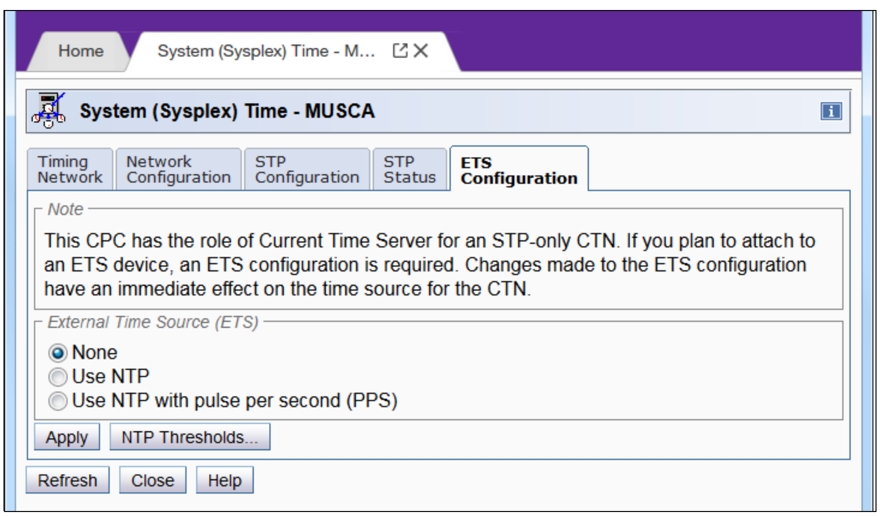

Figure 5-39 System (Sysplex) Time: ETS Configuration

Figure 5-39 shows that None were selected for any ETS.

11. If you click **Use NTP**, you are presented with the window to define an NTP Time server and also Query the status of that NTP Time server.

12. Click **Select**. Then, select **Configured** and enter a valid IP address.

13. Click **Query** to test the defined NTP server definition, as shown in Figure 5-40.

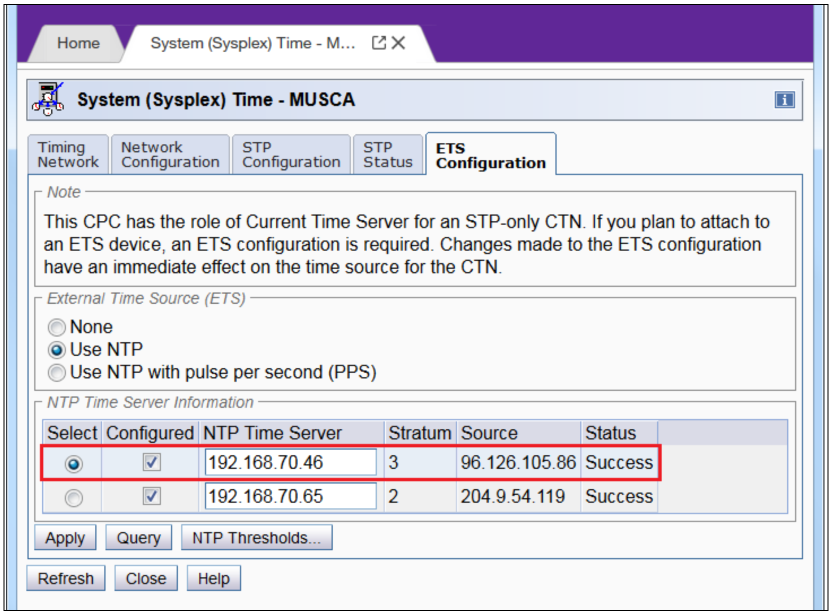

Figure 5-40 System (Sysplex) Time: ETS configuration

If the Query is successful, the value of the Stratum level, Source and Status table fields, and the Status column are automatically completed. The status displays Success if the related NTP server is accessible.

14. Click the **STP Status** tab to observe and verify any CF links that connect to this or other CPCs.

Observe Local STP Link Identifier **0506** that we displayed earlier in the Channel Problem Determination window for CHPID 1.EC (see Figure 5-41).

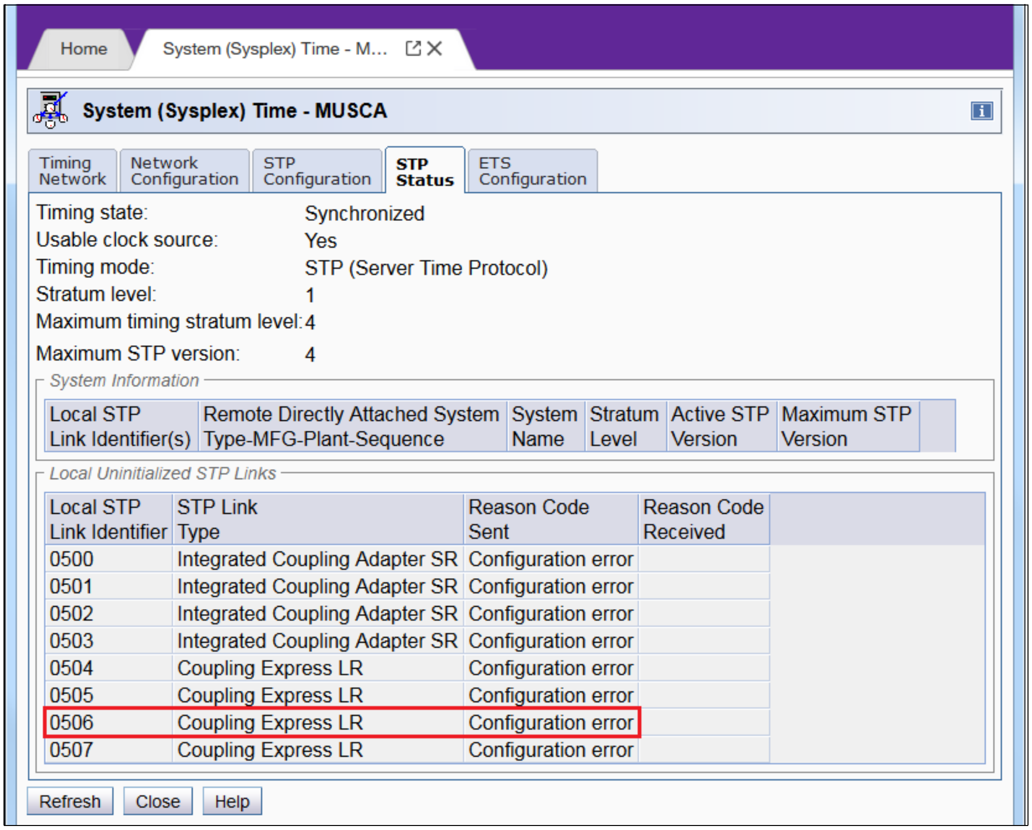

Figure 5-41 System (Sysplex) Time: STP status

In this example, a Configuration error is displayed in the Reason Code Sent field because this server detected a different CTN ID on one of the links to the attached server. However, this situation might not always be the case, depending on your environment.

For more information about how to set up the STP configuration, see Chapter 8, "Preparing for Sysplex and configuring Server Time Protocol" on page 159.

15. Click the **STP Configuration** tab to display or set the CTN ID (STP ID) (see Figure 5-42).

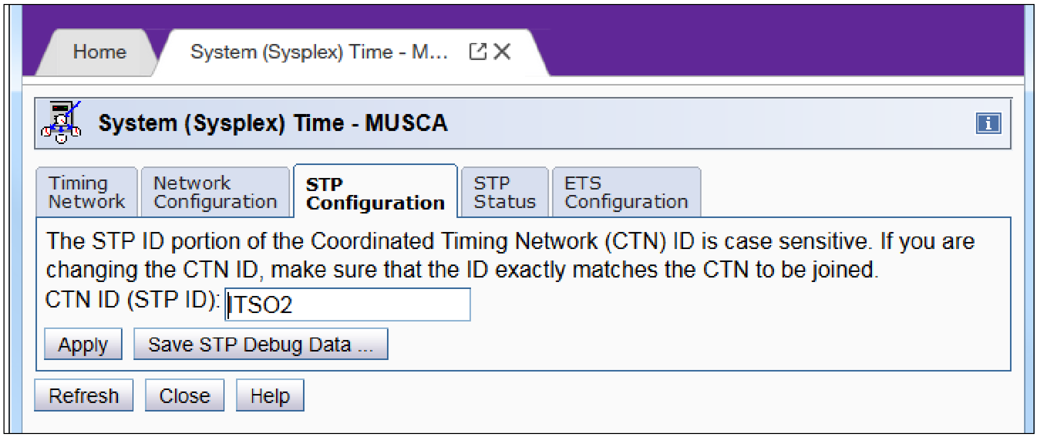

Figure 5-42 System (Sysplex) Time: STP Configuration

16. Click the **Network Configuration** tab to display or set the role for this CPC. In this example, MUSCA is set as the Preferred Time Server (PTS) for the current CTN (see Figure 5-43).

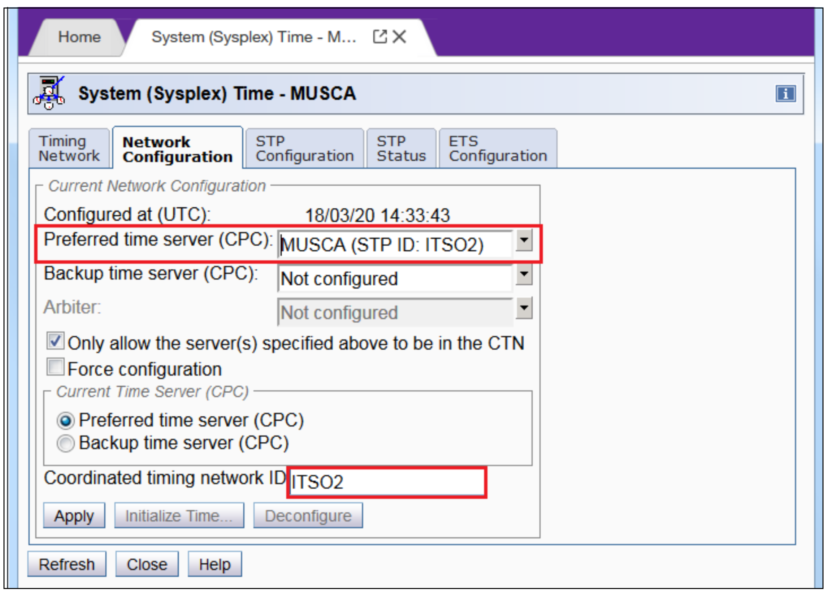

Figure 5-43 System (Sysplex) Time: Network Configuration

17. Click the **Timing Network** tab to check or set the Time for this CPC. Adjustments can also be made here if the CPC lost its ETS and drifted too far from the time zones time (see Figure 5-44).

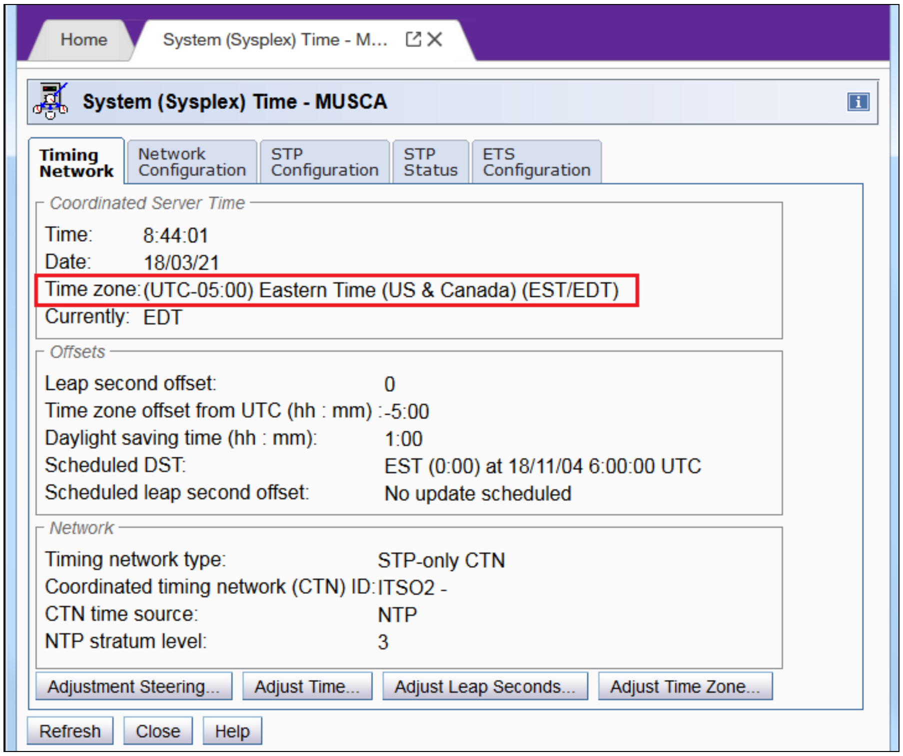

Figure 5-44 System (Sysplex) Time: Timing Network

5.6 Building and verifying Load (IPL) profiles

Now that the CPC is Power on Reset, the Images profiles are defined and activated, the CF links are verified, and the STP and its roles set up, it is time to define a Load (IPL) profile to use to activate (IPL) a partition (LPAR).

To build a Load profile, complete the following steps:

1. Log on by using SYSPROG authority to the HMC for the new 3907.
2. Under Systems Management, click **Systems Management** to expand the list.
3. Under Systems Management, click the radio button that is next to the system to select it (in this example, MUSCA).

4. In the Tasks window, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (see Figure 5-45).

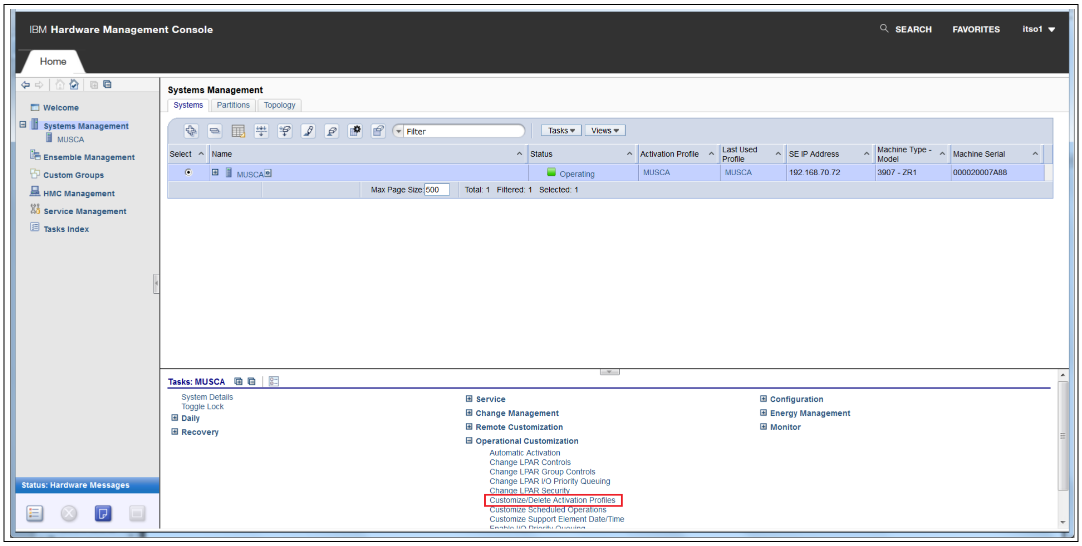

Figure 5-45 Systems Management: Main display

5. Select the DEFAULTLOAD Load Profile and click **Customize profile**.

6. Enter the following required parameters that are specific to your installation to perform an IPL:

   – Profile name: Your preferred profile name over DEFAULTLOAD

   – Description: Your preferred description

   – Load type: Normal

   – Load address: The device address of the IPL volume (97D1)

   – Load parameter: 944301M1:

     - 9443: The device address of the IODF volume
     - 01: The suffix of the LOADxx member in SYS#.IPLPARM on device 9443
     - M: Automatic IPL
     - 1: SYS1.NUCLEUS

Figure 5-46 shows an example.

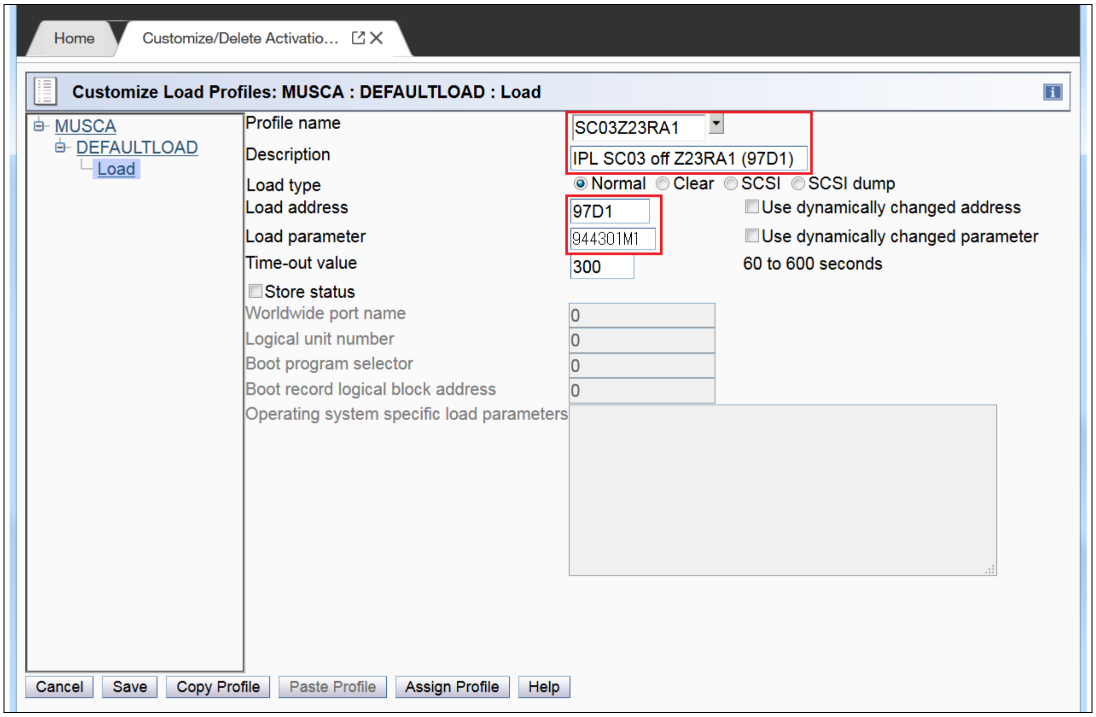

Figure 5-46 Customize Load Profiles: Load

7. Click **Save**. Then, click **OK** to continue to the save window (see Figure 5-47).

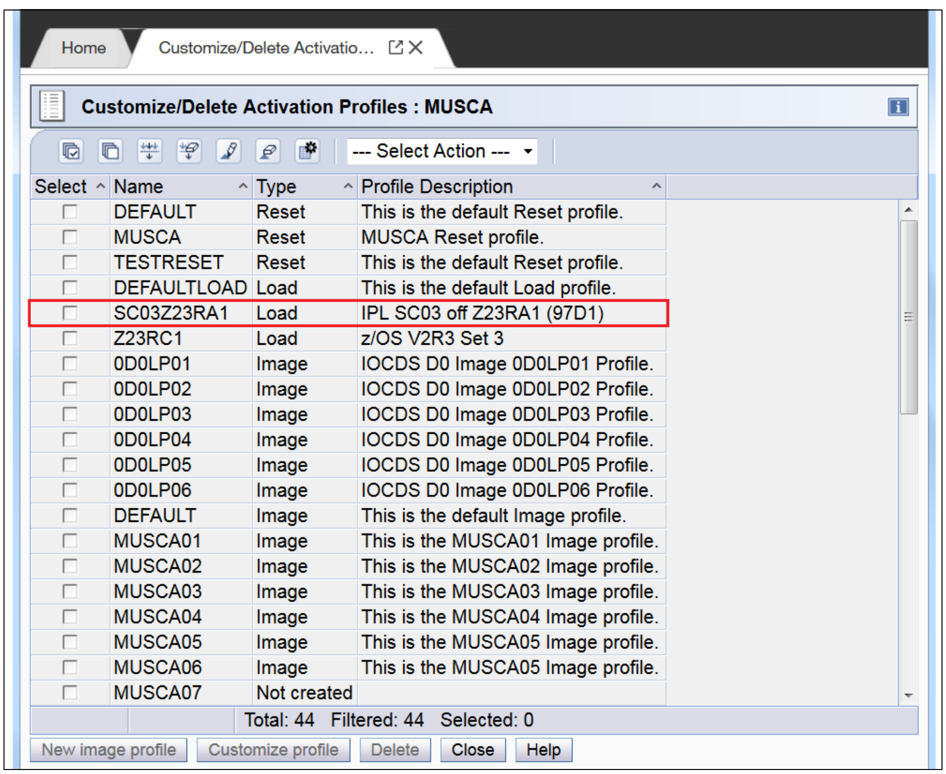

Figure 5-47 Customize Load Profiles: New Load profile

For more information about planning, see Chapter 2, “Planning considerations for CPC in PR/SM mode” on page 7.

5.7 Building and verifying LOADxx Members in SYS#.IPLPARM

As described in 5.6, “Building and verifying Load (IPL) profiles” on page 118, a LOADxx suffix is also required to perform an IPL. This data set member is stored in SYS#.IPLPARM on the volume to which the IODF is written. In our example, this volume is 9443 (IODFPK). The number is the value that you use in your installation for SYS# data sets. This number can be 0 - 9; for example, SYS5.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 MUSCA). Sometimes, the LPARNAME keyword is also used in the LOADxx members and might need to be reviewed or updated, such as MUSCA11.

> **Note:** If you share a LOADxx member with many partitions, the HWNAME and LPARNAME keywords are required.

To build and verify LOADxx Members in SYS#.IPLPARM, complete the following steps:

1. Log on to a system by using TSO that can access 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 member LOADxx. Figure 5-48 on page 122 shows the settings that are used in our example:
   - HWNAME keyword set to MUSCA.
   - LPARNAME keyword set to MUSCA11.
   - IODF keyword set to ** (where ** directs the IPL to look at what IODF/IOCP was PORed into the CPCs HSA, then looks on the IODF volume for that corresponding IODF). A specific IODF suffix number can also be defined in the LOADxx member if you need to override the HSA match.
   - IODF keyword also points to the High-Level Qualifier of the IODF data set (SYS6) and also the 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 NIP consoles and Esoterics.
   - The remaining parameters are use for z/OS and not for the IODF.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
-------------------------------------------------------------------------------
EDIT       SYS0.IPLPARM(LOAD01) - 01.99                   Columns 00001 00072
Command ===>                                                  Scroll ===> CSR
000128 *-----------DEFINITION FOR SC03-------------*
000129 HWNAME   MUSCA
000130 LPARNAME MUSCA11
000131 SYSPLEX  DSTZOS3  Y
000132 IODF     ** SYS6     ITSO     01 Y
000133 SYSCAT   DZ3CAT123CMCAT.SC03
000134 PARMLIB  SYS1.PARMLIB
000135 PARMLIB  SYS1.IBM.PARMLIB
000136 PROCVIEW CORE,CPU_OK
000137 *
000138 HWNAME   CETUS
000139 LPARNAME CETUS26
000140 SYSPLEX  DSTZOS3  Y
000141 IODF     ** SYS6     ITSO     01 Y
000142 SYSCAT   DZ3CAT123CMCAT.SC03
000143 PARMLIB  SYS1.PARMLIB
000144 PARMLIB  SYS1.IBM.PARMLIB
000145 PROCVIEW CORE,CPU_OK
 F1=Help      F2=Split     F3=Exit      F4=Expand    F5=Rfind     F6=Rchange
 F7=Up        F8=Down      F9=Swap     F10=Left     F11=Right    F12=Cancel
```

Figure 5-48 z/OS: SYS#.IPLPARM: LOADxx member

5.8 Communicating information about the new z14 ZR1

Now that you created a 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 network to your z14 ZR1. It includes the following sections:

6.1 Preparation for defining and customizing OSA-Express

To define OSA-Express to your I/O configuration, you need following information:

- OSA operation mode (CHPID type)
- PCHID of OSA-Express and the CHPID number that is associated with that PCHID
- CHPID access list and candidate list within the LCSS
- CNTLINIT number and IODEVICE number

For more information about these operations, see Chapter 13, “Adding network devices” on page 293.

Depending on your network configuration and environment, you can perform OSA-Express customization by using OSA Advanced facilities. OSA advanced facilities are used for customizing the following settings:

- OSA Address Table (OAT) and SNA timer
- Physical port speed
- MAC address

For more information about these settings, see 6.3, “Customizing OSA-Express using OSA Advanced facilities” on page 127.

6.2 Defining OSA Express to your I/O configuration

Defining the OSA-Express involves the following high-level steps:

- Choosing the OSA-Express CHPID type
- Defining the OSA-Express to IO configuration
- Confirm your OSA-Express IO definition

6.2.1 Choosing the OSA-Express CHPID type

Before OSA-Express is used, you must choose the CHPID type. The CHPID types that are supported in z14 ZR1 are listed in Table 6-1. For more information, see the *Open Systems Adapter-Express Customer's Guide and Reference,* SA22-7935, and *IBM Z Connectivity Handbook,* SG24-5444.

Table 6-1 OSA-Express CHPID types that are supported in z14 ZR1

CHPID type	OSA-Express operation mode
OSE	Non-Queued Direct Input/Output mode for SNA and TCP/IP networking
OSD	Queued Direct Input/Output mode for TCP/IP networking
OSC	OSA Integrated Console Controller
OSX	OSA Express for Intraensemble Data Network (IEDN)
OSM	OSA Express for Intraensemble Management Network (INMN)

Note: CHPID type OSN is *not* supported in z14 ZR1.

6.2.2 Defining the OSA-Express to IO configuration

You must define the CHPID, CNTLUNIT, and IODEVICE parameters for using OSA-Express. For more information about how to define I/O configuration using HCD, see Chapter 13, “Adding network devices” on page 293.

6.2.3 Confirm your OSA-Express IO definition

You can confirm your definition by using the following z/OS command:

```
DISPLAY M=CHP(xx)
```

You can confirm the CHPID path status by using the command that is shown in Figure 6-1.

```
D M=CHP(B9)

IEE174I 14.25.14 DISPLAY M 744
CHPID B9:  TYPE=11, DESC=OSA DIRECT EXPRESS, ONLINE
DEVICE STATUS FOR CHANNEL PATH B9
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
01B9 +  +  +  +  +  +  +  +  +  +  +  +  .  .  .  +
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 016C
PNETID 1 = PERFNET
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 6-1 OSD D M=CHP

Confirm the channel path to a device by using the following command:

```
DISPLAY M=DEV(xxxx)
```

The node descriptor information that is returned includes the emulated control unit 1730.008 and 1732.001, which are used for the OSA-Express6S port. Also included is the 3907 machine type and serial number, as shown in Figure 6-2.

```
D M=DEV(1B90)

IEE174I 14.39.13 DISPLAY M 751
DEVICE 01B90   STATUS=ONLINE
CHP                   B9
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             1B90
INTERFACE ID          6C00
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 001730.008.IBM.02.390700007A88.B900
SCP TOKEN NED     = 001730.008.IBM.02.390700007A88.B900
SCP DEVICE NED    = 001732.001.IBM.02.390700007A88.B900
```

Figure 6-2 The OSD M=DEV command

Confirm the device number and the status by using the following command:

```
DISPLAY U
```

The device number and type of device that is defined are shown in Figure 6-3.

```
D U,,,1B90,1

IEE457I 14.42.59 UNIT STATUS 753
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
1B90 OSA  A-BSY                                   0


D U,,,1B9F,1

IEE457I 14.45.21 UNIT STATUS 755
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
1B9F OSAD O-RAL                                   0
```

Figure 6-3 The OSD D U,,,device command

6.3 Customizing OSA-Express using OSA Advanced facilities

OSA Advanced facilities is a tool that is integrated in the HMC. To start OSA Advanced facilities, log in to the HMC with the correct authority, and select the CPC that requires OSA customization. Then, select **Operational Customization** → **OSA Advanced Facilities** (see Figure 6-4).

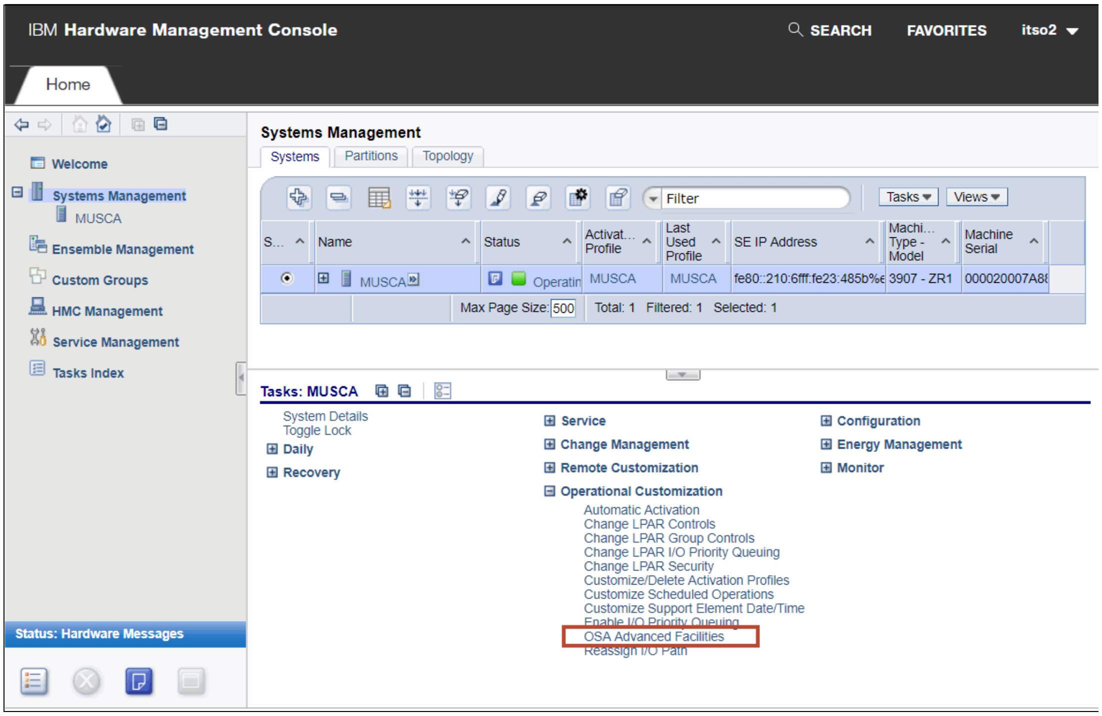

Figure 6-4 HMC OSA Advanced facility selection

The OSA Advanced Facilities window that is shown in Figure 6-5 opens.

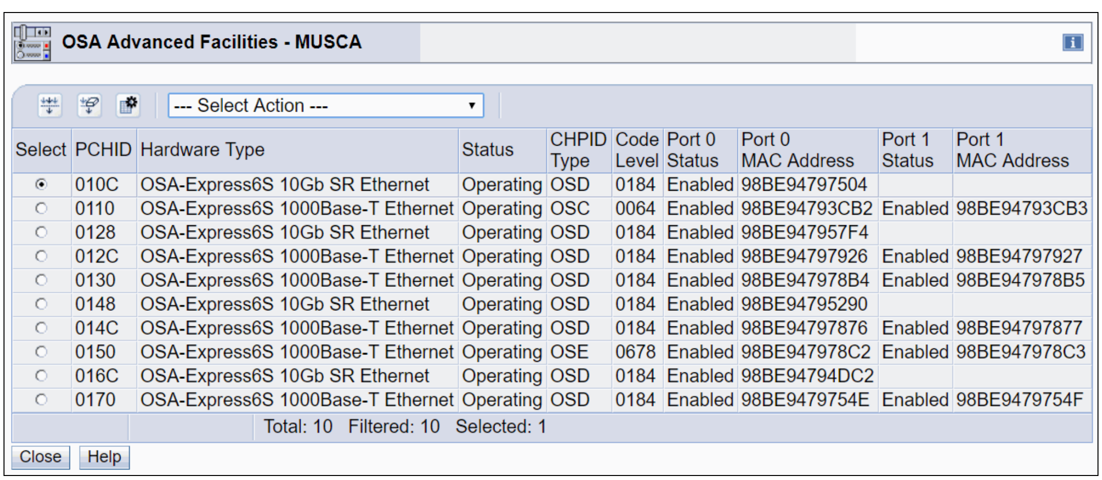

Select	PCHID	Hardware Type	Status	CHPID Type	Code Level	Port 0 Status	Port 0 MAC Address	Port 1 Status	Port 1 MAC Address
⦿	010C	OSA-Express6S 10Gb SR Ethernet	Operating	OSD	0184	Enabled	98BE94797504		
○	0110	OSA-Express6S 1000Base-T Ethernet	Operating	OSC	0064	Enabled	98BE94793CB2	Enabled	98BE94793CB3
○	0128	OSA-Express6S 10Gb SR Ethernet	Operating	OSD	0184	Enabled	98BE947957F4		
○	012C	OSA-Express6S 1000Base-T Ethernet	Operating	OSD	0184	Enabled	98BE94797926	Enabled	98BE94797927
○	0130	OSA-Express6S 1000Base-T Ethernet	Operating	OSD	0184	Enabled	98BE947978B4	Enabled	98BE947978B5
○	0148	OSA-Express6S 10Gb SR Ethernet	Operating	OSD	0184	Enabled	98BE94795290		
○	014C	OSA-Express6S 1000Base-T Ethernet	Operating	OSD	0184	Enabled	98BE94797876	Enabled	98BE94797877
○	0150	OSA-Express6S 1000Base-T Ethernet	Operating	OSE	0678	Enabled	98BE947978C2	Enabled	98BE947978C3
○	016C	OSA-Express6S 10Gb SR Ethernet	Operating	OSD	0184	Enabled	98BE94794DC2		
○	0170	OSA-Express6S 1000Base-T Ethernet	Operating	OSD	0184	Enabled	98BE9479754E	Enabled	98BE9479754F

Total: 10 Filtered: 10 Selected: 1

Figure 6-5 OSA Advanced Facilities initial window

6.3.1 Configuring OAT and SNA LLC2 timer for OSE channel

When you define OSA-Express as OSE channel (non-QDIO mode), you must customize OSA Address Table (OAT) *except* for the following uses:

- Use *only* the default OAT and do *not* use (require) port sharing.
- OSA-Express Direct SNMP subagent.

For more information, see *Open Systems Adapter-Express Customer's Guide and Reference,* SA22-7935.

You can edit OAT with Open Systems Adapter/Support Facility (OSA/SF). OSA/SF is a tool for customizing OSA-Express. From zEC12 GA2, the OSA/SF tool is available on HMC. It is included as a submenu of the OSA Advanced facilities.

You must use OSA/SF on the HMC to define OSA Address Table (OAT) and SNA definition for OSA-Express5S/6S. You can still use OSA/SF running on z/OS for OSA-Express4S, but generally migrate to OSA/SF on HMC.

The OSA/SF tool capabilities are compared in Table 6-2.

Table 6-2 OSA/SF comparison

Type of OSA/SF	Supported OSA Express
On HMC	OSA Express-4S/5S/6S
On z/OS, z/VM	OSA Express-4S

For other CHPID types, OSA/SF is not required. If you need to change the port speed or the Media Access Control (MAC) address, you can configure it from the OSA Advanced Facility. The OSA CHPID type and OSA/SF requirements are listed in Table 6-3.

Table 6-3 OSA/SF and OSA CHPID reference

OSA CHPID type	OSA/SF
OSE	Required
OSD	Not required
OSC	Not supported
OSX	Not supported
OSM	Not supported

Note: OSN channel is not supported on z14.

When you use SNA network with OSE CHPID type and need to change SNA timer (SNA LLC2 parameter), you must customize the parameters in OSA/Advanced facilities. For more information, see *Open Systems Adapter-Express Customer's Guide and Reference,* SA22-7935.

If you are upgrading from z13s to z14 ZR1, and your configuration uses the OSA Express-5S 1000Base-T feature, your OSE configuration is automatically migrated. After this process is complete, check the configuration in OSA/SF on the HMC.

Customizing OAT and SNA timer

> **Note:** Before you customize OAT and SNA timer with OSA/AF on the HMC, see *Open Systems Adapter/Support Facility on the Hardware Management Console,* SC14-7580. This publication is available for download from IBM Resource Link.

To customize OAT and SNA time for OSE channel, complete the following steps:

1. Start OSA Advanced facilities.
2. The OSA Advanced Facilities window opens (see Figure 6-6). Select the OSE channel PCHID that you want to customize. In this example, we selected **PCHID 0174**. Then, select **Select Action** → **Card specific advanced facilities**, as shown in Figure 6-6.

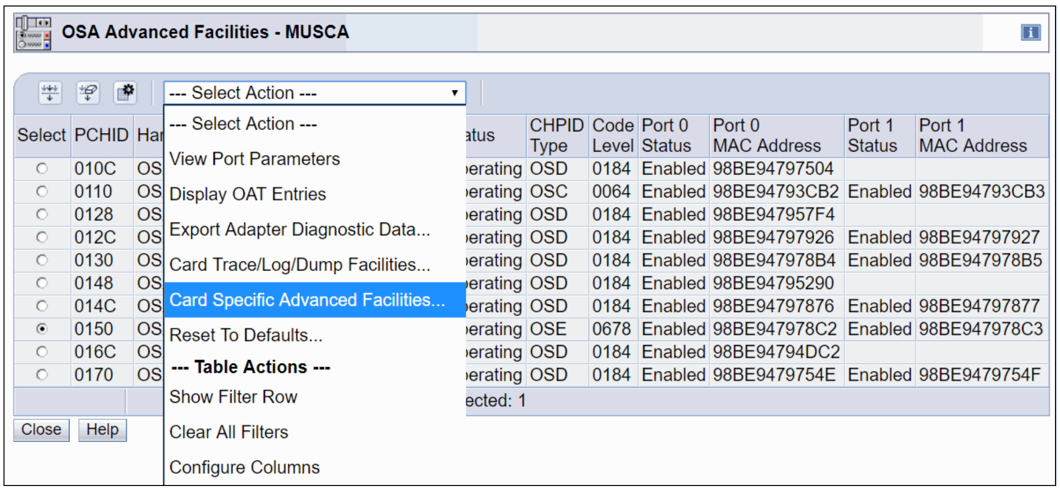

Figure 6-6 Entering OSA/SF on the HMC

3. The Advanced Facilities window opens (see Figure 6-7). To edit OAT and SNA timer entries, select **Panel configuration options** and click **OK**.

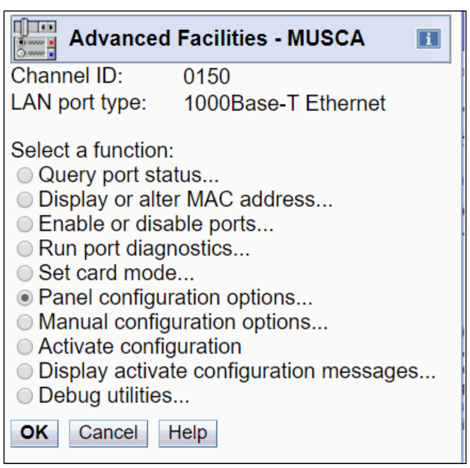

Figure 6-7 OSA/SF on the HMC Panel configuration options

4. The Panel Configuration Options window opens (see Figure 6-8). You can define the following options:

Edit OAT entries By selecting this option, you can edit the OAT and SNA definition. An OAT entry defines the data path between an OSA feature port and an LPAR image.

Edit SNA timers By selecting this option, you can enter SNA timer values.

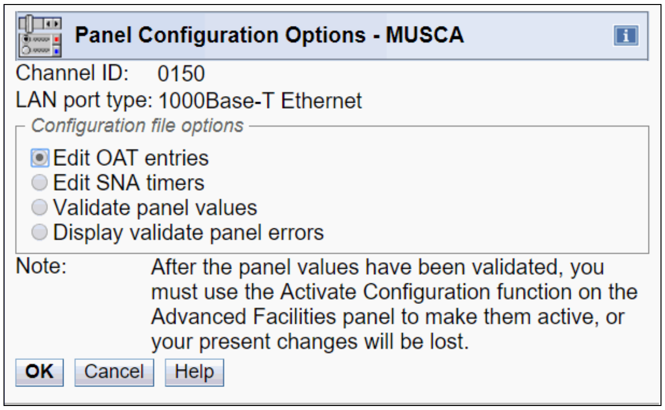

Figure 6-8 OSA/SF on the HMC: Configuration file options

5. Choose **Edit OAT entries**, and the Edit OSA Address Table (OAT) Entries window opens, (see Figure 6-9).

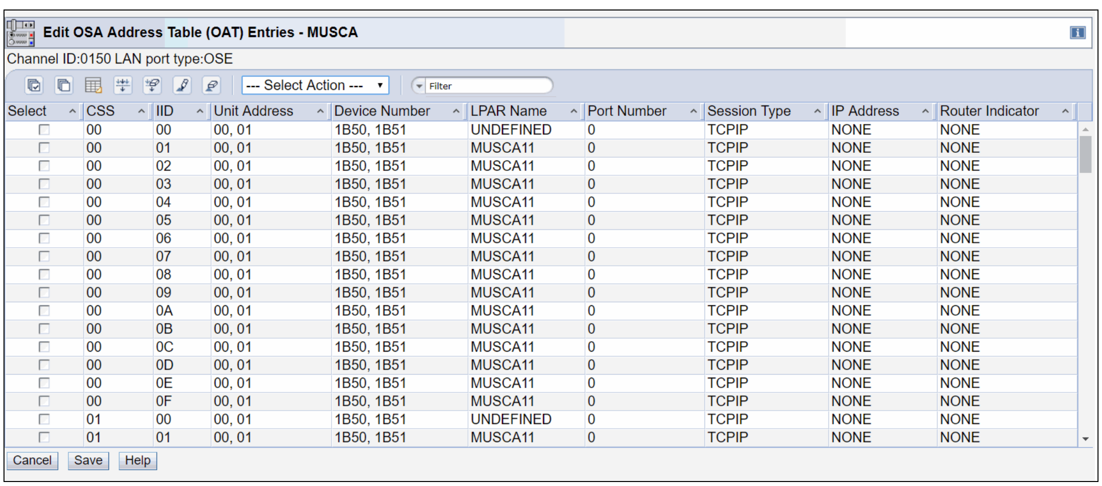

Select	CSS	IID	Unit Address	Device Number	LPAR Name	Port Number	Session Type	IP Address	Router Indicator
	00	00	00, 01	1B50, 1B51	UNDEFINED	0	TCPIP	NONE	NONE
	00	01	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	02	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	03	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	04	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	05	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	06	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	07	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	08	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	09	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0A	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0B	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0C	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0D	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0E	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	00	0F	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE
	01	00	00, 01	1B50, 1B51	UNDEFINED	0	TCPIP	NONE	NONE
	01	01	00, 01	1B50, 1B51	MUSCA11	0	TCPIP	NONE	NONE

Figure 6-9 OSA/SF on HMC: Edit OSA Address Table (OAT) Entries window

6. To edit, select the device in the left column, as shown in Figure 6-10. Click **Select Action** and choose **Edit as TCP/IP entry** or **Edit as SNA entry** for the selected device.

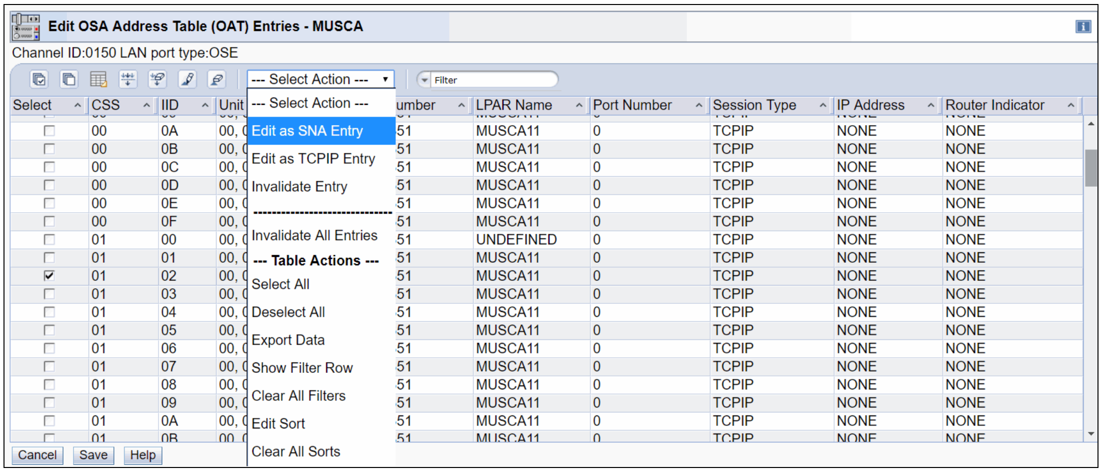

Figure 6-10 OSA/SF on the HMC: Select Action for Edit OSA Address Table (OAT) Entries window

7. Select **Edit as SNA Entry**, and the OSA Address Table (OAT) Entry window that is shown in Figure 6-11 opens. Select the correct SNA entry and click **OK**.

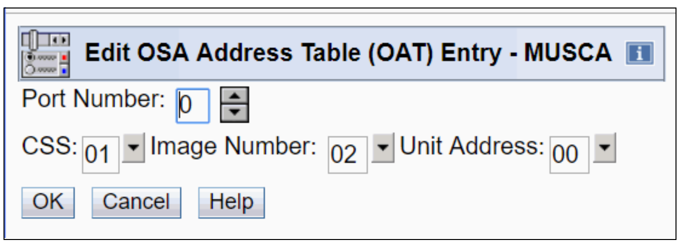

Figure 6-11 OSA/SF on the HMC: Edit SNA entry window

8. Select **Edit as TCP/IP Entry**, and the Edit OSA Address Table (OAT) Entry that is shown in Figure 6-12 opens. Enter and select the appropriate TCP/IP parameters here. Then, click **OK**.

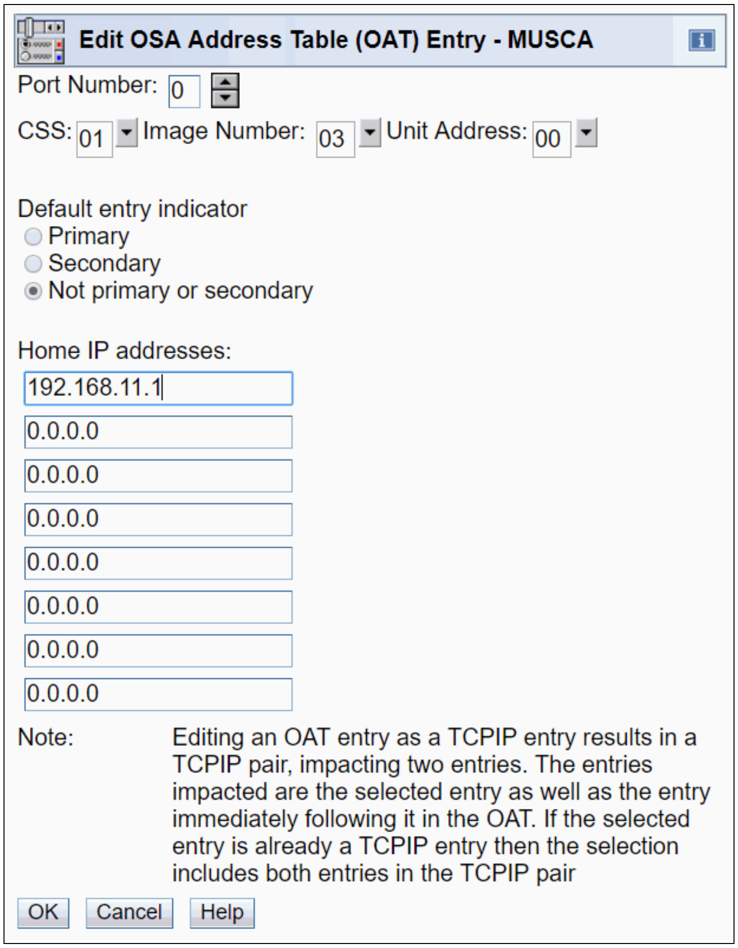

Figure 6-12 OSA/SF on the HMC: Edit TCP/IP entry window

9. After editing TCP/IP entry or SNA entry, the Edit OSA Address Table (OAT) Entries window opens again. Confirm that your entries are displayed in this window. Figure 6-13 shows that the TCP/IP entry and SNA entry can be confirmed. Click **Save** to save the configuration.

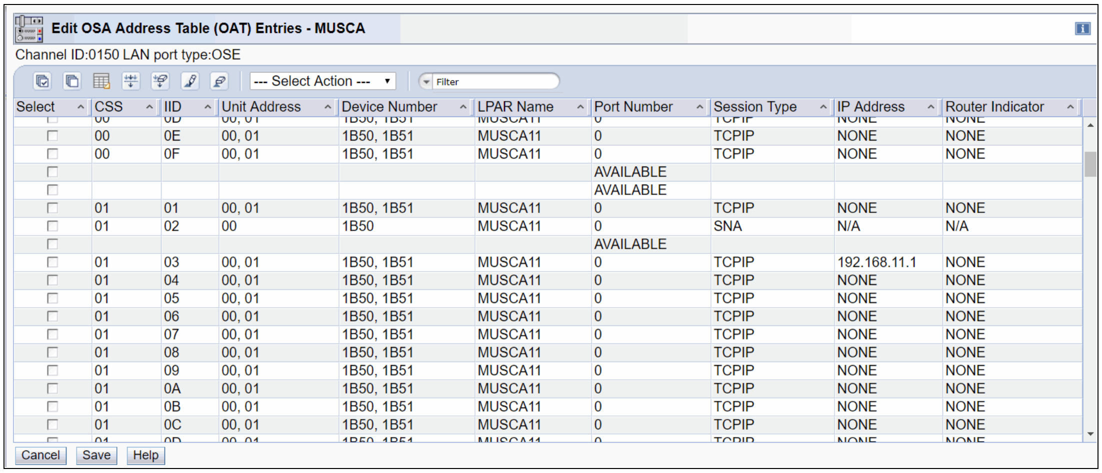

Figure 6-13 OSA/SF on the HMC: Edit OSA Address Table (OAT) Entries window after editing

10. Select **Edit SNA timers** in the Panel Configuration Options window to change the SNA timer setting. Figure 6-14 shows the window in which you can edit the values. You can set the parameters for port 0 and 1 individually.

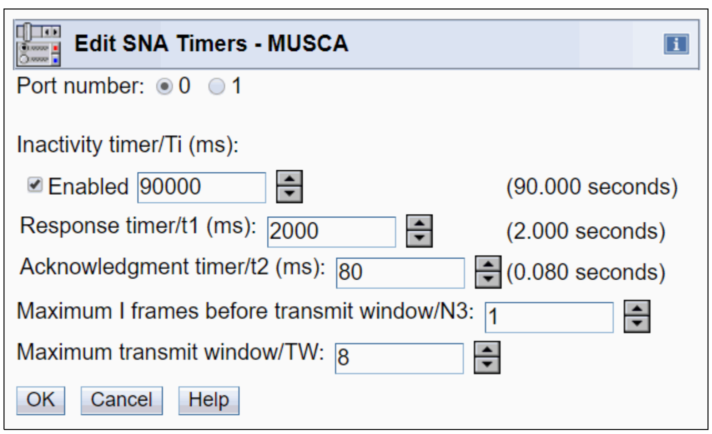

Figure 6-14 OSA/SF on the HMC: Edit SNA Timers window

11. When you complete editing the OAT entries or SNA timer, the Panel Configuration Options window is displayed again. To activate the settings, you must validate them by selecting **Validate panel values**. Then, click **OK** (see Figure 6-15).

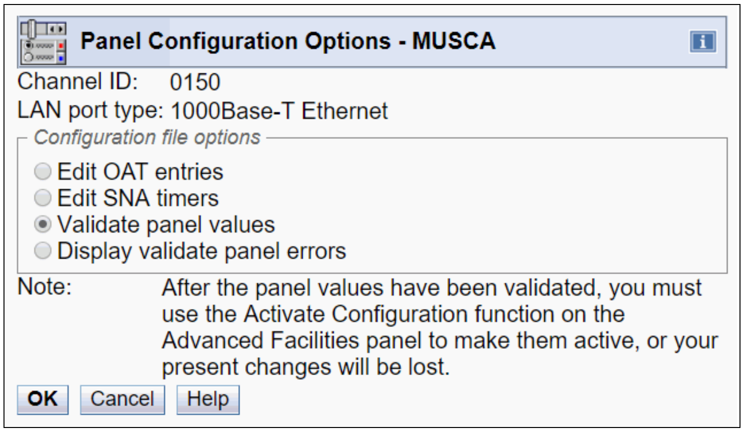

Figure 6-15 OSA/SF on the HMC: Validate panel values

12. If all of the parameters are entered correctly, a window as shown in Figure 6-16 is displayed.

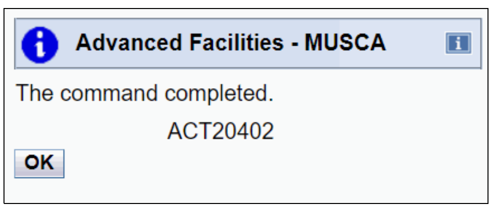

Figure 6-16 OSA/SF on HMC: Validate window value with success

If any of the entered parameters are incorrect, a window with ACT20425 message opens. To identify the error, select **Display validate panel errors**, and correct the error. Then, validate the panel values again until the error is fixed.

13. To activate your OSA configuration, select **Activate configuration** from the Advanced Facilities window, and click **OK** (see Figure 6-17).

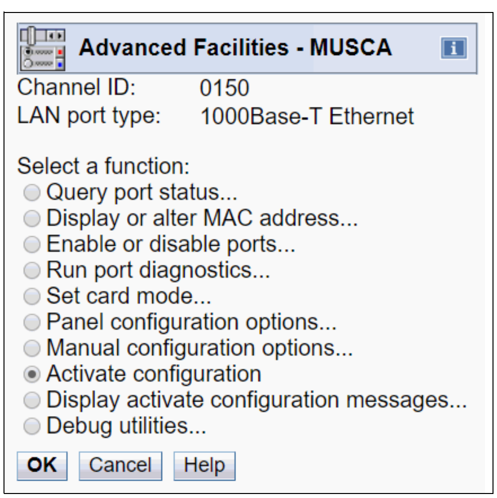

Figure 6-17 OSA/SF on the HMC: Activate configuration

14. The confirmation window that is shown in Figure 6-18 opens. Click **Yes** to continue.

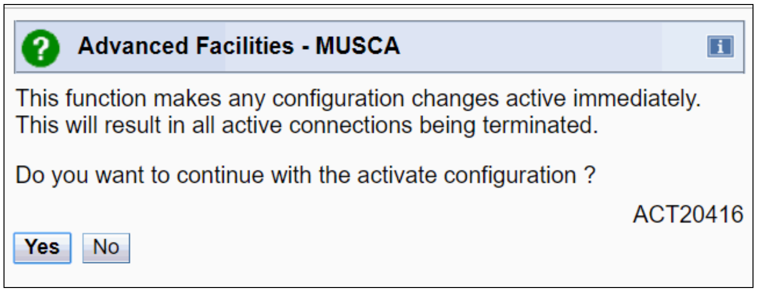

Figure 6-18 OSA/SF on the HMC: Confirm activation

15. When the activation is successfully, the window that is shown in Figure 6-19 is displayed. Click **OK** to complete the process.

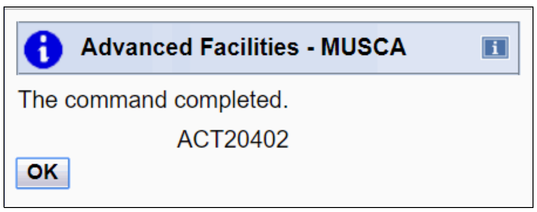

Figure 6-19 OSA/SF on the HMC: message when activation completed

16. To effective the changes, you must bring the CHPID OFFLINE from all the LPARs that share the OSA CHPID, and then back ONLINE.

17. If you select **Manual configuration options** (see Figure 6-7 on page 129), the window that is shown in Figure 6-20 opens. In this window, you can import or export the source file of OAT by using a USB device or FTP, create a configuration file in the editor on HMC, and edit the source file. For more information, see *Open Systems Adapter/Support Facility on the Hardware Management Console,* SC14-7580.

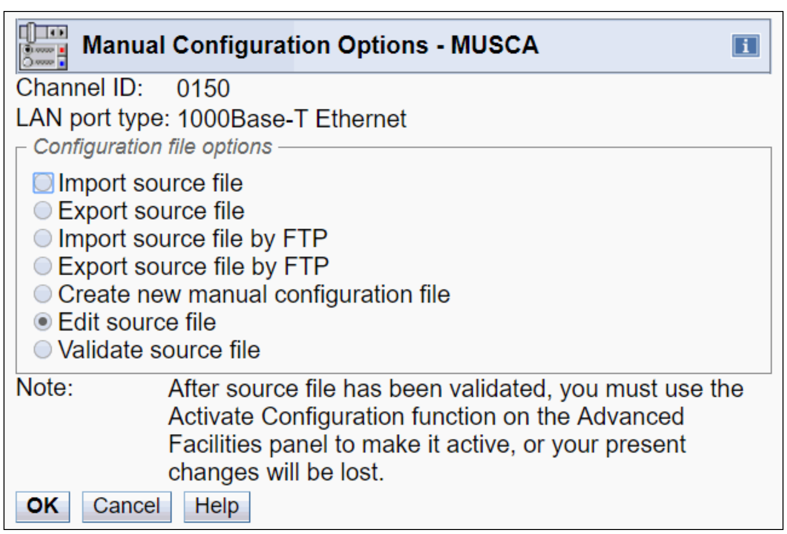

Figure 6-20 OSA/SF on the HMC: Manual Configuration options

6.3.2 Setting OSA parameters using OSA Advanced Facilities

If you want to change the port speed or MAC address of an OSA-Express feature, you can use the OSA Advanced Facilities. Normally, changing MAC address is required for SNA networks.

> **Port speed:** You can change only the port speed to OSA-Express5S/6S 1000BASE-T. Other OSA features do not include this capability.
>
> You cannot set port speed to 1000 Mbps for OSA-Express5S/6Sress4S 1000BASE-T. If you want to set the port speed to 1000 Mbps, you must select **Auto Negotiate**.

Setting the OSA port speed

To change port speed, complete the following steps:

1. Log on to the HMC, select the CPC you want to operate, and then, click **Open OSA Advanced Facilities**.
2. Select the PCHID of the OSA channel for which you need to set the card mode. Select **Card specific advanced facilities** and click **OK**.
3. Select **Set card mode** and click **OK** (see Figure 6-21).

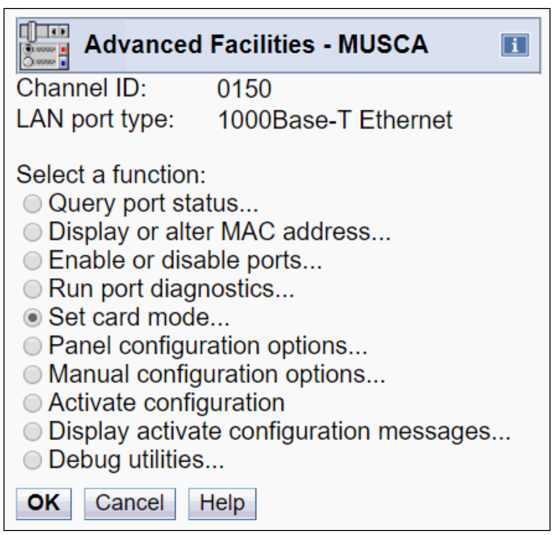

Figure 6-21 Selecting Card Mode definition

4. The Set Card Mode or Speed window opens (see Figure 6-22). Select the correct port speed from the Speed/Mode list. You can set the speed of port 0 and 1 individually. Click **OK**.

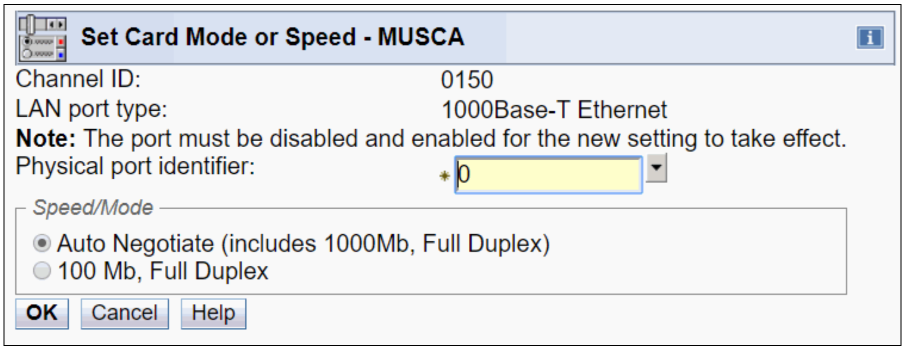

Figure 6-22 Set card mode or speed

5. To effect the changes, you must configure CHPID OFFLINE and ONLINE from every LPAR where this CHPID is defined.

Changing OSA MAC address

To change the MAC address, complete the following steps:

1. Log on to the HMC.
2. Open OSA Advanced Facilities, and select the PCHID that you want to customize.
3. The Advanced Facilities window opens (see Figure 6-23). Select **Display or alter MAC address** and then, click **OK**.

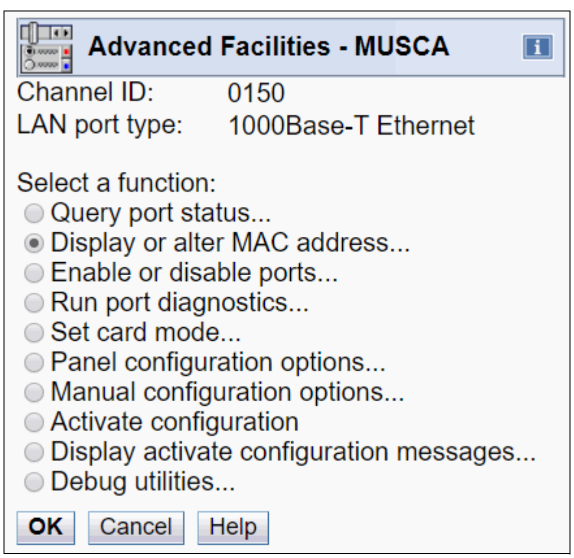

Figure 6-23 Display or alter MAC address

4. The Display or alter MAC address window opens (see Figure 6-24). Set the MAC address that you want and then, click **OK**.

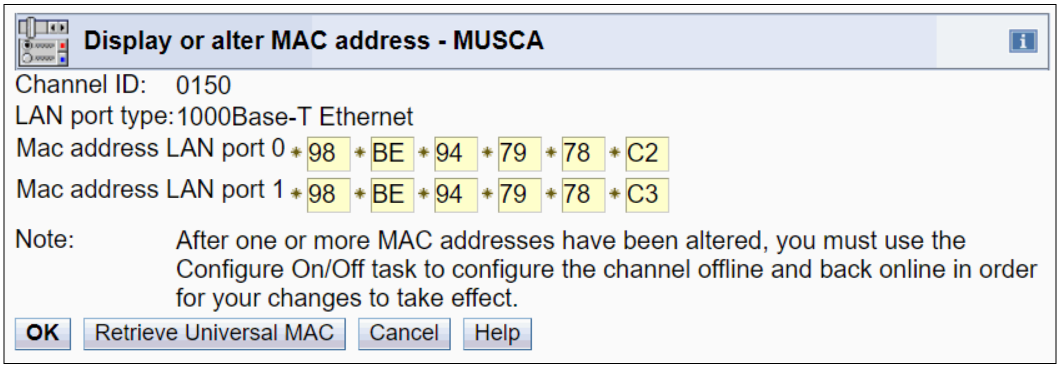

Figure 6-24 Display or alter MAC address values

5. To reflect the modification, you must configure CHPID OFFLINE and ONLINE from every LPAR where this CHPID is defined.

6.3.3 Confirming your OSA customization

To confirm your customization on the OSA Advanced Facility, open OSA Advanced Facility again and confirm whether your definitions are reflected correctly. Also, check the device status from your networking software, such as TCP/IP and VTAM.

6.4 Shared Memory Communications (SMC-R and SMC-D)

For more information about how to configure Shared Memory Communications over RDMA (SMC-R), see 10.4, “Shared Memory Communications over RDMA (SMC-R)” on page 257.

For more information about how to configure Shared Memory Communications - Direct Memory Access (SMC-D), see 10.5, “Shared Memory Communications - Direct Memory Access” on page 261.

6.5 Channel-To-Channel connection

This section describes the configuration steps for channel-to-channel connections (CTC).

6.5.1 FICON CTC: Preparing

If your system requires CTC connection for programs, such as VTAM or GRS, FICON CTC (FCTC) links must be configured. Because ESCON channels are not supported on zEC12 and later IBM Z servers, the only available option is to implement FICON CTC.

The FICON CTC does not require that you explicitly define the CTC control unit function to one channel side or the other. The channel decides which side contains the control unit function. Therefore, both sides of an FCTC connection use an FC channel path that is defined to one or more FCTC control units that represent the target systems for the CTC connections.

A logical address is required when the FC channel path is attached to a shared FC channel path to identify the logical partition with which to communicate. When attached to an unshared FC channel path, the logical address must be zero or not specified. For more information, see the CTC Definitions for z Systems web page.

FICON CTC communication within a single processor complex can be accomplished with a single FC channel path that is connected to an FC switch. That is, both sides of the CTC connection are represented by the same single FC channel path. This configuration results in the following advantages:

- Reducing the number of channels required
- Simplified configuration design and definition

6.5.2 FICON CTC: Implementation

To define FICON CTC, the following considerations apply to all FICON CTC configurations:

- The server at each end of an FICON CTC connection uses a FICON native (CHPID type FC) channel.
- The FICON native channel at each end of the CTC connection features a defined FICON CTC control unit.
- The FICON CTC devices on the FCTC control unit are defined as type FCTC.
- The FCTC control function on the IBM Z platform can communicate with an FCTC control unit that is defined on a FICON native channel on any server that supports FICON.
- The FICON native channel at each end of the FICON CTC connection supports the FCTC control units. It also communicates with other FICON native control units, such as disk and tape.

In a FICON CTC configuration, FICON CTC control units are defined at each end, but only one end provides the FICON CTC control unit function. During initialization of the logical connection between two ends of a FICON CTC connection, the channel that provides the FICON CTC control unit function is determined by using an algorithm. This process results in balancing the number of FCTC CU functions that each end of the logical connection is providing. The algorithm uses the channel with the lower FC worldwide name (WWN) to provide the FICON CTC control unit function.

The FICON native channel CTC communication does *not* require a pair of channels because it can communicate with any FICON native channel that includes a defined corresponding FCTC control unit. This configuration means that FICON CTC communications can be provided by using only a single FICON native channel per server.

For more information about how to implement FICON CTC, see the following publications:

- *FICON CTC Implementation*, REDP-0158
- *FICON Planning and Implementation Guide*, SG24-6497
- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804

6.5.3 FICON CTC: Management

After activating the new FICON CTC configuration and all cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the Support Element (SE) windows:

- Checking status by using z/OS commands. If you are interested in checking the status of CHPIDs 44 and 45, issue **D M=CHP(44)** and **D M=CHP(45)** commands, as shown in Figure 6-25.

```
D M=CHP(44)
IEE174I 16.42.20 DISPLAY M 081
CHPID 44:  TYPE=1B, DESC=FICON SWITCHED, ONLINE
DEVICE STATUS FOR CHANNEL PATH 44
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
04A1 +  +  +  +  .  .  .  .  .  .  .  .  .  .  .  .
SWITCH DEVICE NUMBER = 0001
DEFINED ENTRY SWITCH - LOGICAL SWITCH ID = 01
ATTACHED ND = 008960.F64.IBM.CA.1000010546MH
PHYSICAL CHANNEL ID = 0175
FACILITIES SUPPORTED = ZHPF
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    § PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL

D M=CHP(45)
IEE174I 16.43.04 DISPLAY M 083
CHPID 45:  TYPE=1B, DESC=FICON SWITCHED, ONLINE
DEVICE STATUS FOR CHANNEL PATH 45
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
04A1 .  .  .  .  .  .  .  .  +  +  +  +  .  .  .  .
SWITCH DEVICE NUMBER = 0002
DEFINED ENTRY SWITCH - LOGICAL SWITCH ID = 02
ATTACHED ND = 008960.F64.IBM.CA.1000010546MD
PHYSICAL CHANNEL ID = 0115
FACILITIES SUPPORTED = ZHPF
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    § PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 6-25 DISPLAY FCTC on system SC03

- Checking status by using SE windows:
  - From the HMC, select the CEC (under Systems Management) from where the CHPID/PCHID you want to verify is stored and click **Single Object Operations** (under the **Recovery** task options).
  - On the Support Element, select the same CEC and click **Channels**. Look for the PCHID of which you are interested in checking the status. The result is shown in Figure 6-26.

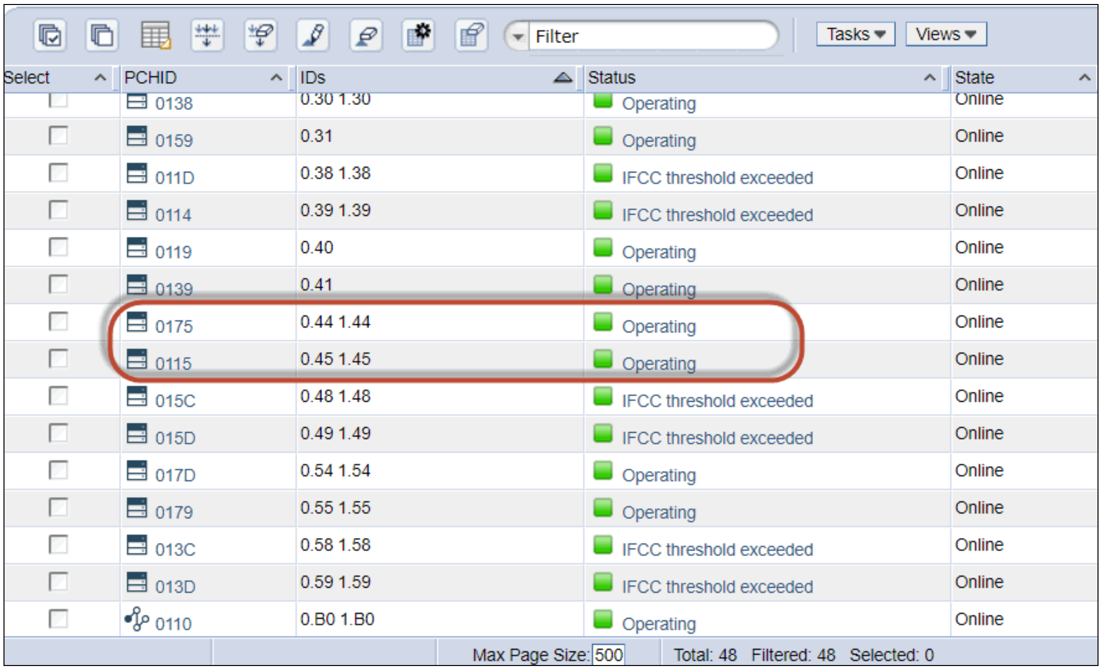

Figure 6-26 Using the SE to verify channel FCTC using the CEC view

  - For more information about the PCHID, click the PCHID to open the details window, see shown in Figure 6-27.

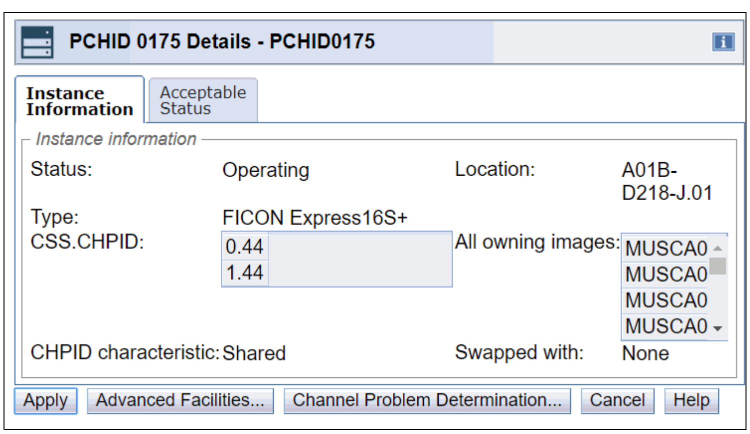

Figure 6-27 FCTC PCHID details

– The status also can be checked 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 as shown in Figure 6-28.

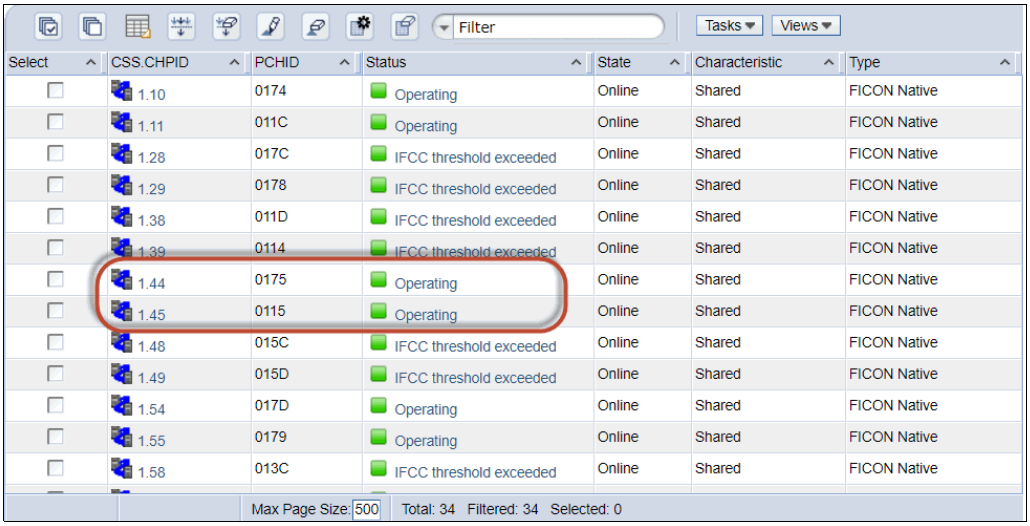

Figure 6-28 Verifying channel FCTC using LPAR view

– For more information, click the CHPID, as shown in Figure 6-29.

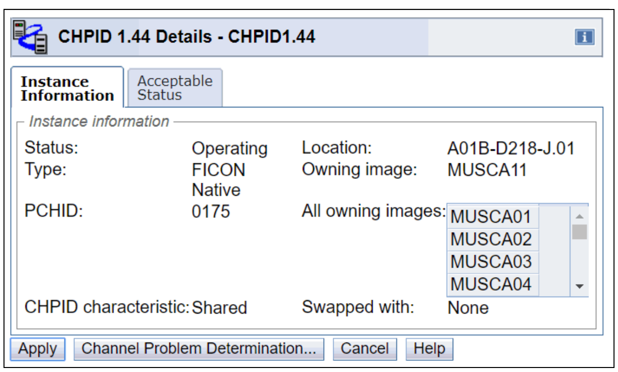

Figure 6-29 FCTC CHPID details

7

Defining console communication

This chapter includes the following topics:

7.1 Preparing console definition

A non-SNA console is a console that is required for IPL, z/OS system operation, and management. You must define at least one non-SNA console to your z/OS system. To define the non-SNA console, use the OSA-Express Integrated Console Controller (OSA-ICC) function. For more information, see *z Systems Open Systems Adapter Integrated Console Controller User's Guide*, SC27-9003.

The process includes the following steps:

1. Define the OSC CHPID and the CNTLUNIT.
2. Define the 3270-X IODEVICE.
3. Configure OSA-ICC by using the OSA Advanced facilities.
4. Export or import OSA-ICC configuration (optional).
5. Activate OSA-ICC configuration by using the OSA Advanced facilities.
6. Setup Personal Communications.

Note: The OSC channel can be defined only for OSA-Express5s/6s 1000Base-T cards.

You must configure OSA-ICC when you upgrade to z14 ZR1 from an z13s server or when a new z14 ZR1 is installed.

For an upgrade from z13s, the OSA-ICC configuration is automatically upgraded to z14 ZR1.

7.2 Defining the OSA-ICC

Before you configure OSA-ICC, you must define the OSC CHPID, CNTLUNIT, and 3270-X IODEVICE to the I/O configuration.

For more information about defining OSC channel by using HCD, see Chapter 13, “Adding network devices” on page 293, and Chapter 11, “Adding LPARs and OSCONFIGs” on page 269.

7.3 Defining a new OSA-ICC configuration by using OSA Advanced facilities

When a new z14 ZR1 is installed, you must configure OSA-ICC from scratch. To define a new OSA-ICC configuration, complete the following steps:

1. Log on to the HMC, select the CPC, and open OSA Advanced facilities.
2. Select the OSC PCHID to be used for the new OSA-ICC configuration, and select **Card specific advanced facilities** (see Figure 7-1).

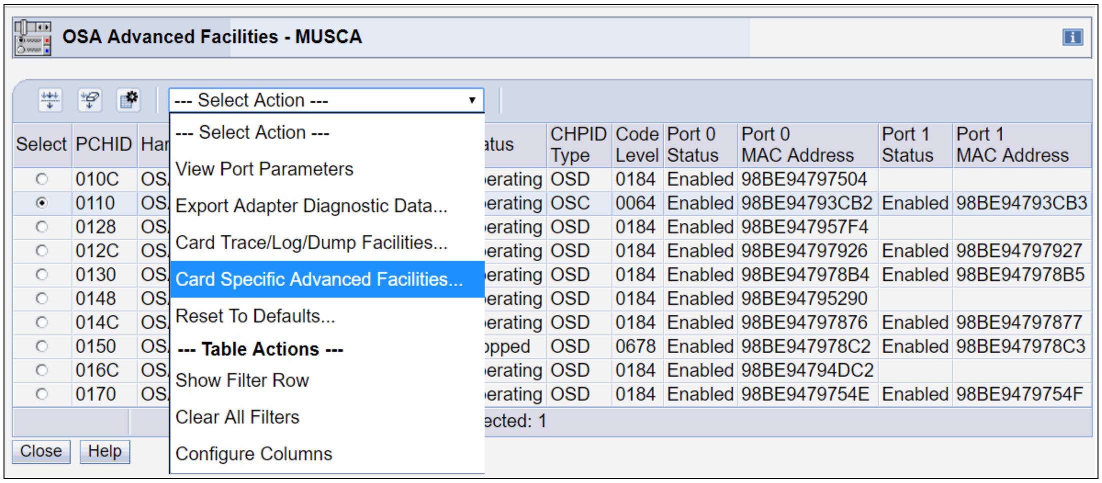

Figure 7-1 Card Specific Advanced facilities

3. Select **Panel configuration options** and click **OK** (see Figure 7-2).

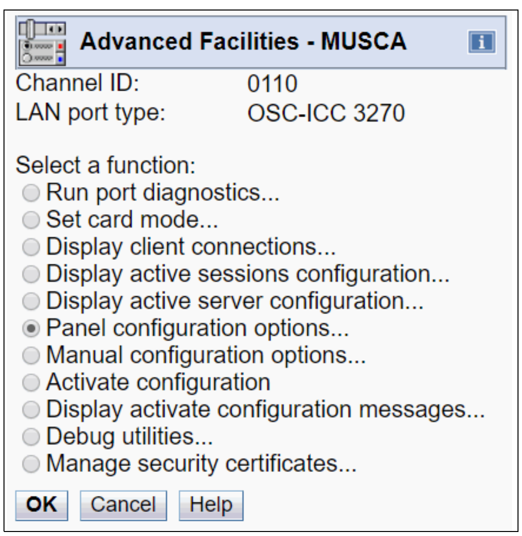

Figure 7-2 Selecting Panel configuration options option

4. The Panel Configuration Options window opens (see Figure 7-3). Define the session and server configurations in this window, and then, validate those values. In this example, we define the server configuration first. Next, we define the session configuration. To edit the server configuration, select **Edit server configuration** and click **OK**.

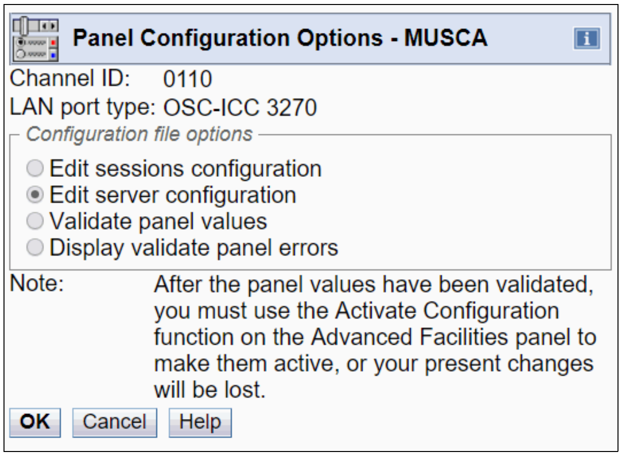

Figure 7-3 Selecting Edit server configuration option

5. The Edit Server Configuration panel opens. Enter the necessary values on this panel. Figure 7-4 shows our sample configuration. Click **OK** to save.

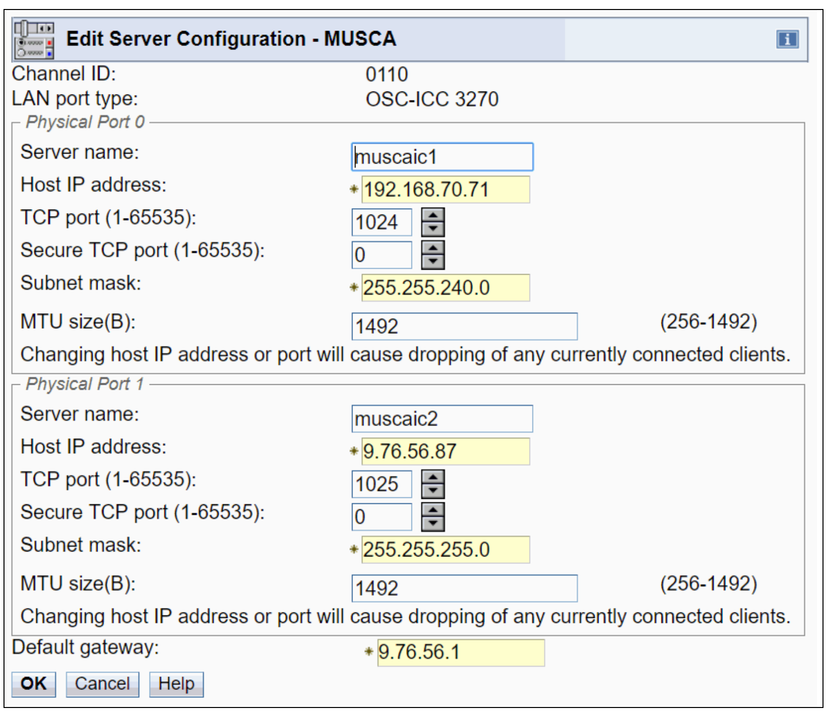

Figure 7-4 Edit Server Configuration panel

6. The command is completed (ACT20402) window opens. Click **OK** to return to the Panel Configuration Options window.

7. Define the session configuration. Select **Edit sessions configuration** and click **OK**. The window that is shown in Figure 7-5 opens.

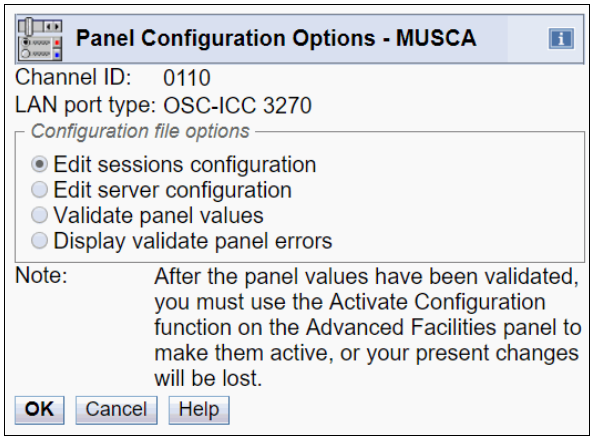

Figure 7-5 Selecting Edit sessions configuration option

8. The Edit Sessions Configuration window opens (see Figure 7-6). To configure a new session, select a number from the Session Index column and click **Change**.

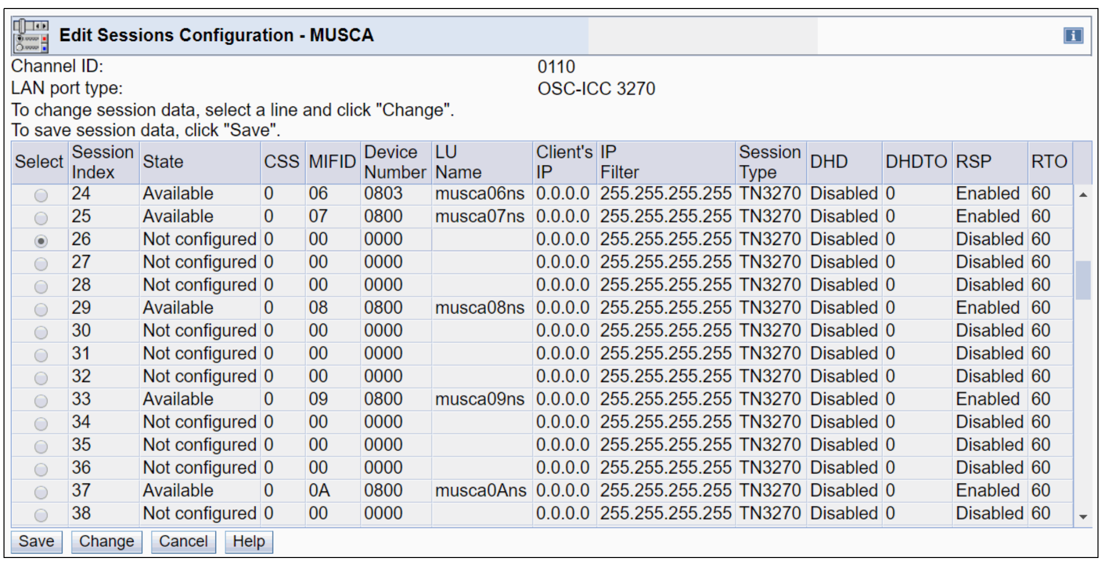

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client's IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
	24	Available	0	06	0803	musca06ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
	25	Available	0	07	0800	musca07ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
	26	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	27	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	28	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	29	Available	0	08	0800	musca08ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
	30	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	31	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	32	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	33	Available	0	09	0800	musca09ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
	34	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	35	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	36	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	37	Available	0	0A	0800	musca0Ans	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
	38	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Figure 7-6 Edit Sessions Configuration selection

9. The Edit Session Configuration panel opens. Define the session parameter here. Click **OK** to save (see Figure 7-7).

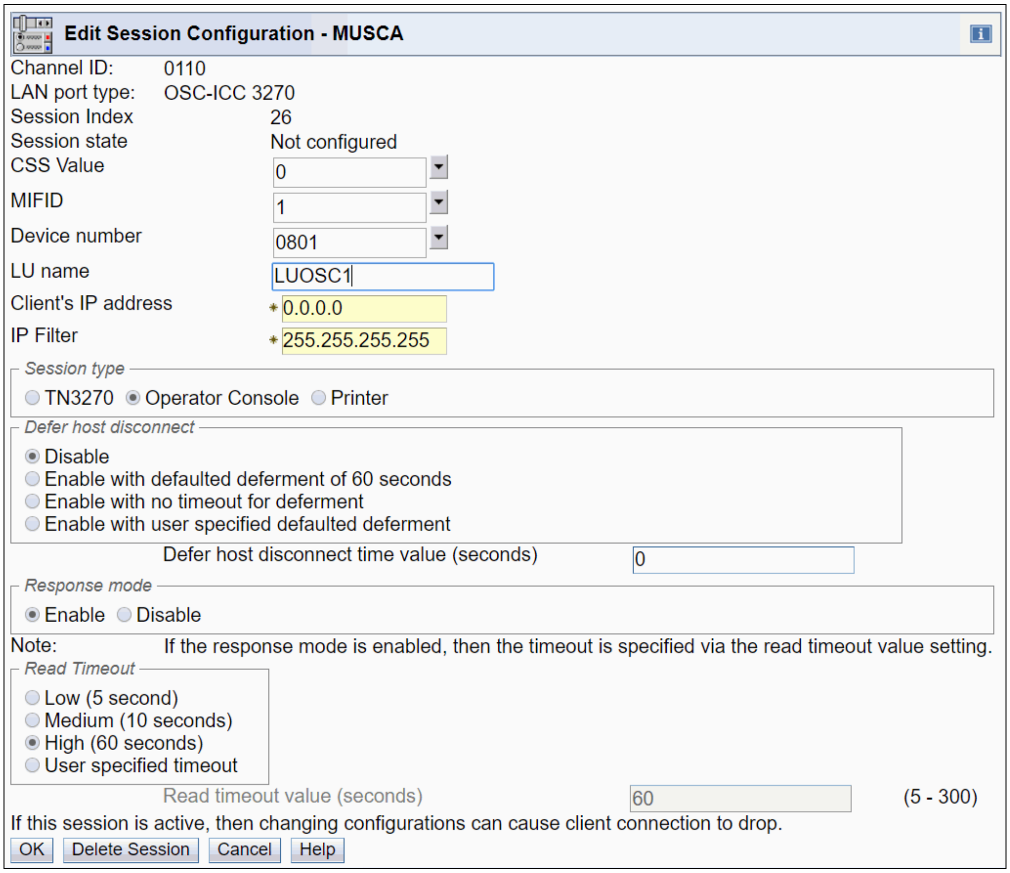

Figure 7-7 Edit Session Configuration input

10. The Edit Sessions Configuration panel is displayed again (see Figure 7-8). Be sure that your input values are displayed correctly. To save session values, click **Save**.

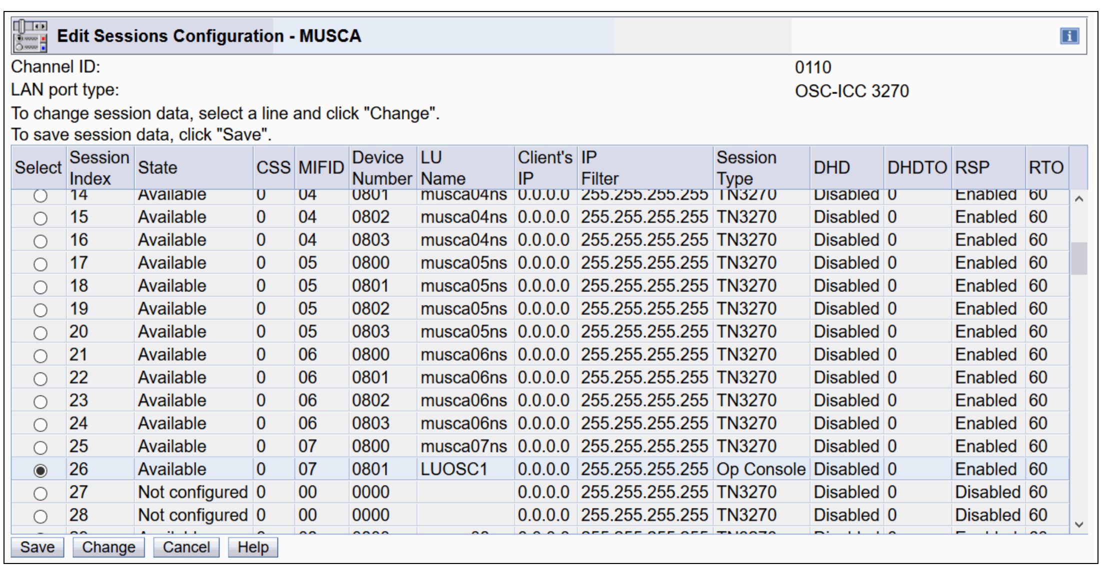

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client's IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
○	14	Available	0	04	0801	musca04ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	15	Available	0	04	0802	musca04ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	16	Available	0	04	0803	musca04ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	17	Available	0	05	0800	musca05ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	18	Available	0	05	0801	musca05ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	19	Available	0	05	0802	musca05ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	20	Available	0	05	0803	musca05ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	21	Available	0	06	0800	musca06ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	22	Available	0	06	0801	musca06ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	23	Available	0	06	0802	musca06ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	24	Available	0	06	0803	musca06ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	25	Available	0	07	0800	musca07ns	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
◉	26	Available	0	07	0801	LUOSC1	0.0.0.0	255.255.255.255	Op Console	Disabled	0	Enabled	60
○	27	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	28	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Figure 7-8 Edit Sessions Configuration after define values

11. The command completed (ACT20402) window opens. Click **OK**.

12. The Panel Configuration Options window opens again (see Figure 7-9). Now, you can validate the values. Select **Validate panel values** and click **OK**.

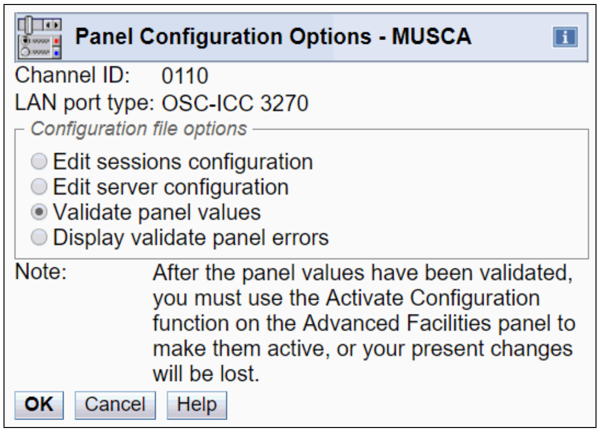

Figure 7-9 Display validate panel

13. If configuration data is correct, the command completed (ACT20402) window opens.

    If an error is encountered, a window similar to the window that is shown in Figure 7-10 opens. You can confirm the error by selecting **Display validate panel errors** on Panel Configuration Options (see Figure 7-9). Correct the error, and select **again the Validate panel values** to check again.

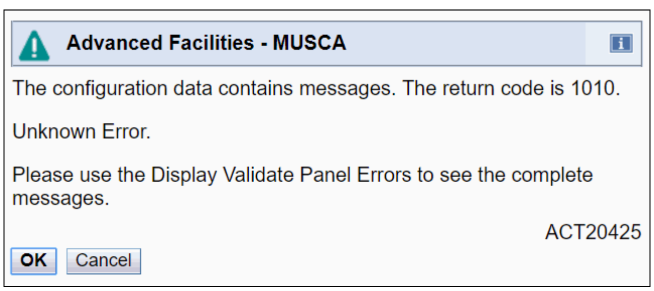

Figure 7-10 Panel Configuration Options error found

14. When the validation is complete without any errors, you can activate the OSA-ICC configuration. To activate configuration, exit the **Panel Configuration Options** window by clicking **Cancel**. Then, select **Activate configuration** in the Advanced Facilities window (see Figure 7-11).

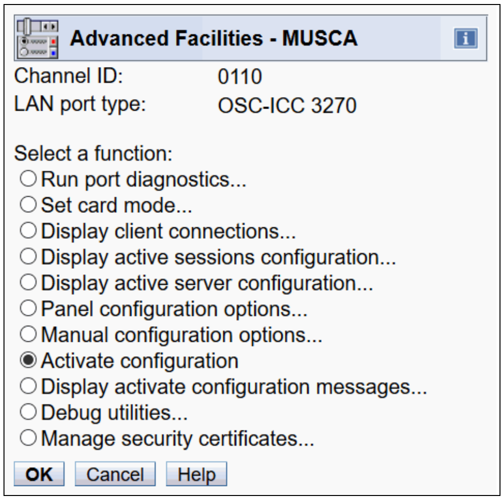

Figure 7-11 Activate configuration

15. The confirmation panel appears. Click **Yes** to continue (see Figure 7-12).

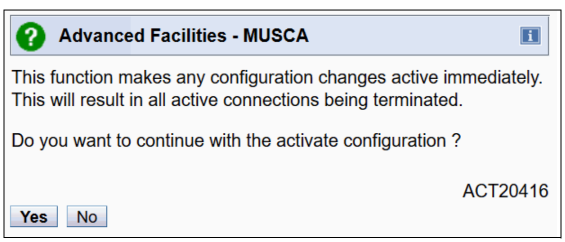

Figure 7-12 Activate configuration confirmation

16. When the activation is complete, the message shown in Figure 7-13 is displayed.

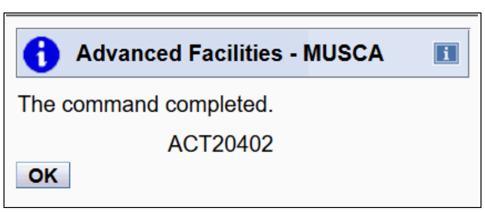

Figure 7-13 Activate configuration success

7.3.1 Saving and restoring OSA-ICC configuration

During an upgrade, you can export an OSA-ICC configuration file from a z13s and import it to the z14 ZR1. This section describes how to export and import the OSA-ICC configuration file by using the HMC.

Exporting OSA-ICC configuration file using OSA Advanced Facilities

In this example, we export the OSA-ICC configuration file from a z13s server and import the file to the (new) z14 ZR1 with a USB flash drive. You can also import or export the configuration file by using FTP.

To export the OSA-ICC configuration file, complete the following steps:

1. Before you export a source file, prepare the USB flash drive that is supported by the HMC, and insert it into the USB port of the HMC. When the flash drive is recognized by the HMC, the window that is shown in Figure 7-14 is displayed.

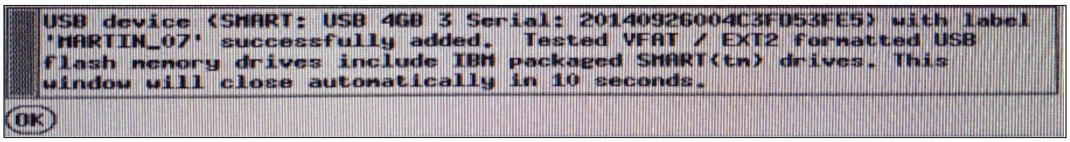

Figure 7-14 Removable media inserted

2. Log on to the HMC, select the CPC you want to operate, and open the OSA Advanced facility.
3. Select the OSC CHPID to export the OSA-ICC configuration file. Next, select **Card Specific Advanced Facilities.** Select **Manual configuration options** and click **OK** (see Figure 7-15).

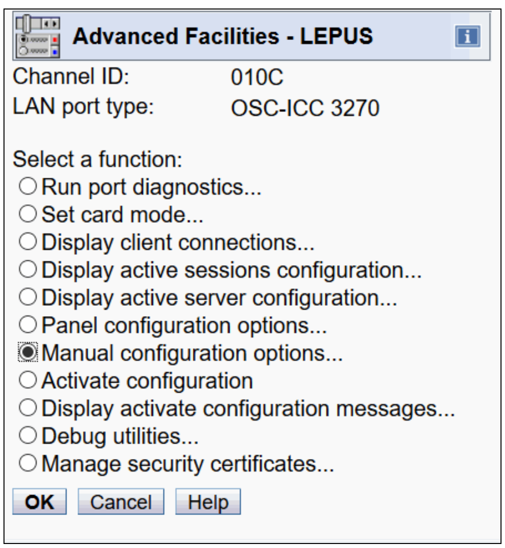

Figure 7-15 Manual configuration options

4. The Manual Configuration Options window opens (see Figure 7-16). Select **Export source file** and click **OK**.

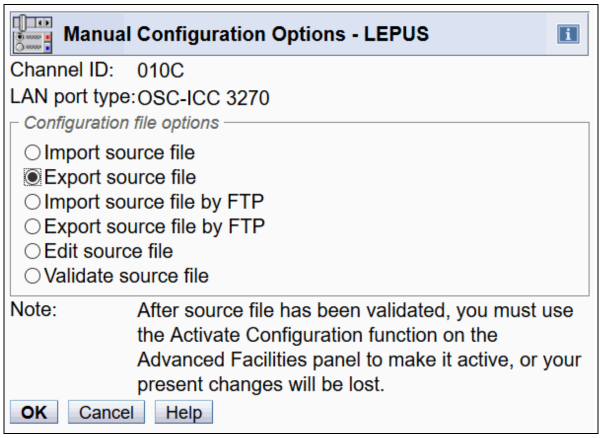

Figure 7-16 Export source file

5. The task requests that a name for the file is written to the inserted media device. For our example, we enter `OSC010C.txt` (see Figure 7-17). Then, click **OK**.

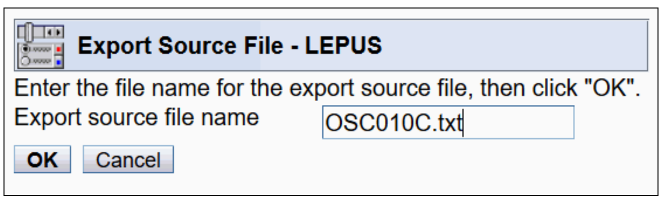

Figure 7-17 Export file: Specify file name

6. The HMC displays the ACT20421 window (see Figure 7-18). Click **OK**.

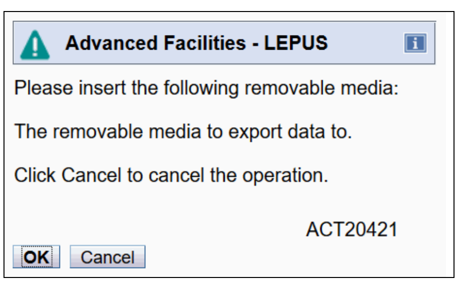

Figure 7-18 ACT20421 window display

7. The HMC task writes the source file for the PCHID that was selected onto the media device and displays a message when it completes (see Figure 7-19). Click **OK**.

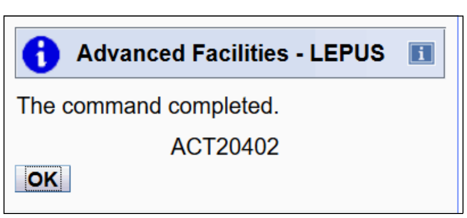

Figure 7-19 Export source file completed

8. You can now remove the USB flash drive. When disconnected, the HMC message advises you that the USB flash drive was removed (see Figure 7-20). Click **OK**.

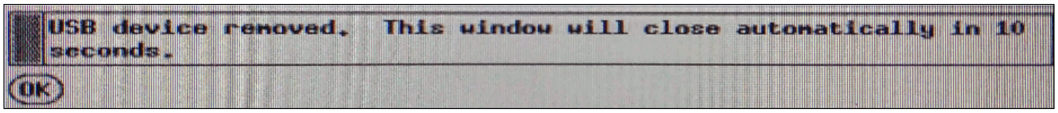

Figure 7-20 Removable media disconnected from HMC

9. Click **Cancel** to exit all of the OSA Advanced Facilities windows.

Example 7-1 shows a source file on the USB flash drive.

Example 7-1 OSA-ICC - sample source file

```
<OSC_SERVER>
<OSC_PHYSICAL_PORT0>
  HOST_IP= 192.168.70.32
  SUBNET_MASK= 255.255.255.0
  PORT= 1024
  SECURE_PORT= 0
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= lepusic1
</OSC_PHYSICAL_PORT0>

<OSC_PHYSICAL_PORT1>
  HOST_IP= 9.76.56.76
  SUBNET_MASK= 255.255.255.0
  PORT= 1025
  SECURE_PORT= 0
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= lepusic2
</OSC_PHYSICAL_PORT1>

DEFAULT_GATEWAY= 9.76.56.1
</OSC_SERVER>

<CONFIG_SESSION>
<SESSION1>
  CSS= 00 IID= 01 DEVICE= 0100
  GROUP=  "lepus01oc"
  CONSOLE_TYPE= 2    RESPONSE= OFF    READ_TIMEOUT= 60
</SESSION1>
...
...
</CONFIG_SESSION>
```

Editing the source file for OSA-ICC

When your OSA-ICC configuration for z14 ZR1 (such as IODEVICE, CSSID, and MIFID of 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 z14

To import the source file and activate the configuration of the OSC-ICC on the new system (z14 ZR1), complete the following steps:

1. Before you import the source file, insert the USB flash drive that contains OSA-ICC source file. When the USB flash drive is recognized by the HMC, the window that is shown in Figure 7-14 on page 151 opens.
2. Log on to the HMC, select the CPC 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** and click **OK** (see Figure 7-21).

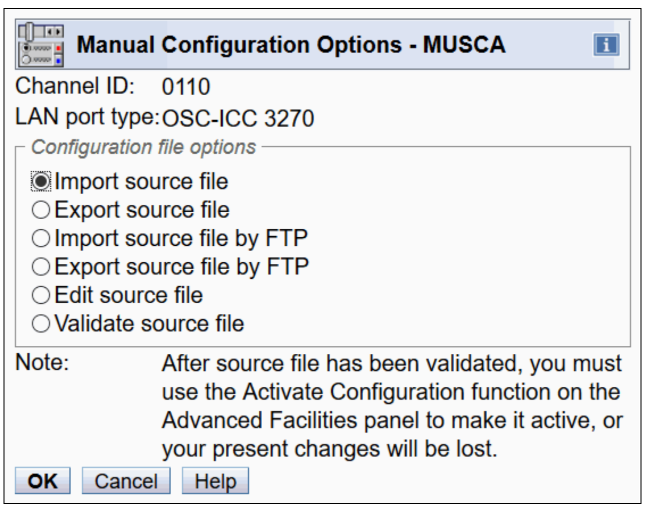

Figure 7-21 Importing a source file

6. You are prompted to insert the media (see Figure 7-22). Click **OK**.

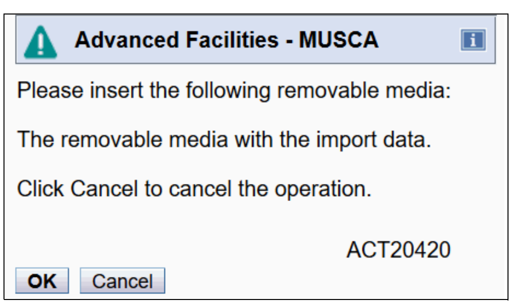

Figure 7-22 Insert media

7. The Import Source File window opens. Select the name of the source file that you want to import. In our example, we select **OSC010C.txt** as the source file (see Figure 7-23) and click **OK**.

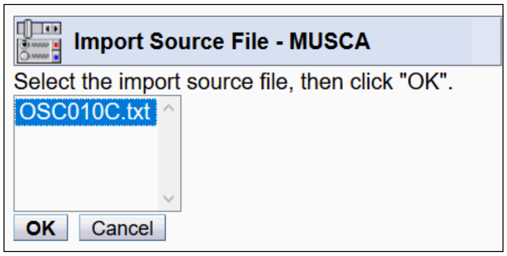

Figure 7-23 Import source file selection

8. The next window (see Figure 7-24) indicates that the source file import is complete. Click **OK** to continue. Then, remove the USB flash drive.

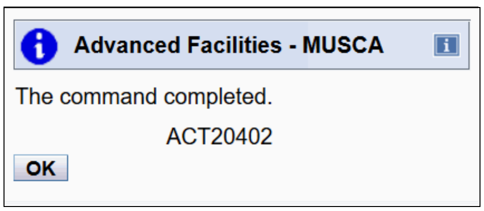

Figure 7-24 Import file successful

After the source file is imported, you must validate it and activate the configuration.

7.4 Verifying the OSA-ICC definition

This section describes some tools that you can use to verify the OSA-ICC 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 (see Figure 7-25).

```
D M=CHP(B0)
IEE174I 16.35.14 DISPLAY M 576
CHPID B0:  TYPE=14, DESC=OSA CONSOLE, ONLINE
DEVICE STATUS FOR CHANNEL PATH B0
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
0080 +  $§ $§ $§ $§ $§ $§ $§ $§ $§ $§ $§ $§ $§ $§ $§
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 0110
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE    § PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE   $ PATH NOT OPERATIONAL
```

Figure 7-25 OSC D M=CHP

- `DISPLAY M=DEV(xxxx)`

  By using this command, you can confirm the channel path to a device. The node descriptor information that is returned includes the emulated control unit 2074. Also included is the 3906 machine type and serial number (see Figure 7-26).

```
D M=DEV(0800)
IEE174I 16.37.52 DISPLAY M 580
DEVICE 00800   STATUS=ONLINE
CHP                   B0
ENTRY LINK ADDRESS    ..
DEST LINK ADDRESS     0D
PATH ONLINE           Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL      Y
MANAGED               N
CU NUMBER             1B00
INTERFACE ID          1000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = NOT AVAILABLE
SCP TOKEN NED     = 002074.   .IBM.02.390700007A88.B000
SCP DEVICE NED    = 002074.002.IBM.02.390700007A88.B000
```

Figure 7-26 OSC D M=DEV(xxx)

- `DISPLAY U`

  By using this command, you can confirm the device number and the status. Check whether the device number and type are defined correctly (see Figure 7-27).

```
D U,,,0800,1
IEE457I 16.39.50 UNIT STATUS 582
UNIT TYPE STATUS        VOLSER     VOLSTATE      SS
0800 3277 O                                       0
```

Figure 7-27 OSC D U,,,device

7.4.2 OSA-ICC console initial window

When you complete setting up the IBM Personal Communications, you see the OSA-ICC initial window, as shown in Figure 7-28. Check whether the OSA-ICC definitions are reflected correctly.

```
 **      OSC Index 61 connected to muscaic2 via IP Addr 9.76.56.87:1025      **
 **       LT Index=33 CSSID=01 MIFID=01 CU=0  UA=00  LUName=muscal1ns        **
 **       Type=3907-ZR1 Mfg=IBM SN=000000007A88 CHPID=B0 Status=Active       **



9.57.215.170:60138
```

Figure 7-28 OSA-ICC initial window

If you do not see the initial window, check your definition in the OSA-ICC or IBM Personal Communications session setup.

8

Preparing for Sysplex and configuring Server Time Protocol

This chapter describes the preparation tasks for IBM Parallel Sysplex and how to configure time synchronization with Server Time Protocol (STP).

This chapter includes the following topics:

8.1 Preparing for Parallel Sysplex

If your z14 ZR1 is a member of a Parallel Sysplex or if you create a Parallel Sysplex that involves the IBM z14 ZR1, time synchronization among CPCs is required. For time synchronization, you must use the Server Time Protocol feature (FC #1021). In addition to the time synchronization, the following tasks are required in preparation for Sysplex:

1. Defining CF (coupling) links.
2. Defining FICON CTC connections (optional).
3. Defining the CF LPAR.
4. Defining the CF LPAR image profile in HMC.
5. Configuring STP.

For more information about defining CF (coupling) links, see Chapter 9, “Defining Coupling Facility links” on page 199.

For more information about defining the image profile, see 5.4, “Creating an Image Profile on the 3907 Support Element” on page 99.

8.2 Preparing for non-sysplex system time synchronization

If your z14 ZR1 is not part of a Parallel Sysplex but you want to synchronize the time among multiple CPCs, you need the STP function and STP-only (timing) links. The following configuration steps are required:

1. Defining STP-only (timing) links.
2. Configuring STP.

For more information about defining STP-only (timing) links, see 9.4, “Defining an STP timing-only link by using ICA SR” on page 211.

Note: This chapter describes how to set up STP on a z14 ZR1 server.

8.3 Server Time Protocol overview

STP provides time synchronization among multiple CPCs in a Coordinated Timing Network (CTN). CTN is a collection of servers that are synchronized to a time value that is named 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.

For more information about STP concepts and planning information, see *Server Time Protocol Planning Guide*, SG24-7280.

With z14 ZR1, a new STP stratum level (stratum 4) is supported. Timekeeping information is transmitted over coupling links. A CTN is shown in Figure 8-1.

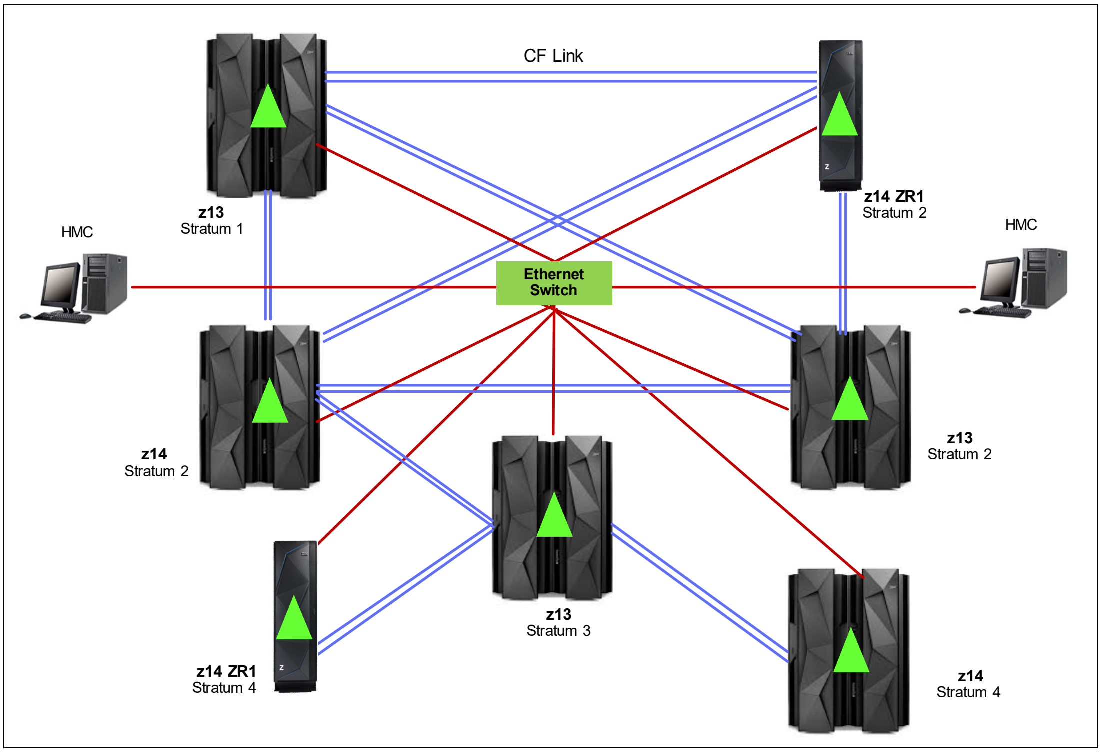

Figure 8-1 STP only CTN with z14 connectivity

Note: z14 ZR1 supports STP timing mode only, which can be part of only an STP-only CTN. The z14 ZR1 cannot be in the same sysplex or CTN with any IBM Z server earlier than a z13.

8.3.1 Using external time source

STP retrieves the time from a Network Time Protocol (NTP) server through the NTP client on the Support Element (SE). STP synchronizes the TOD to the external time source (NTP server) accessed through the NTP client on the SE. To access the ETS, the SE requires network connectivity to the NTP server.

To provide NTP data to the NTP client on the SE, the following options are available:

- Connect NTP server directly to the SE-HMC network.
- Use HMC as NTP server.

Note: A phone modem is no longer supported on the IBM Z. Therefore, you cannot set up the HMC as ETS by using a dial-out configuration. You must provide network connectivity to an NTP server.

For security reasons, do not attach the NTP server directly to the SE-HMC network. HMC can act as an NTP server, so in general use HMC as NTP server for STP. For more information, see 8.4, “Configuring the HMC as an NTP server” on page 162.

A sample ETS configuration for the STP is shown in Figure 8-2.

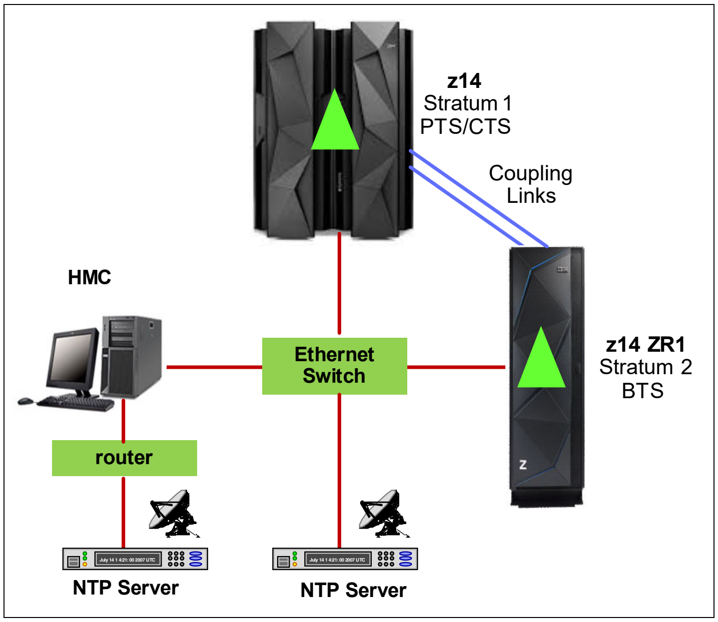

Figure 8-2 Example configuration: NTP server on the HMC

Especially for the financial markets, tight time accuracy is demanded by the authorities of various countries. Therefore, the US Financial Industry Regulatory Authority (FINRA) announced that computer clocks that are used to record events in national market system (NMS) securities and over-the-counter (OTC) equity securities must be synchronized to within a 50-millisecond drift tolerance of the National Institute of Standards and Technology (NIST) atomic clock. Also, the European Union demands in their Markets in Financial Instruments Directive (MIFID) II regulation that the maximum divergence from UTC is to be 100 microseconds.

Unfortunately, the accuracy of the interface with an NTP server to maintain Coordinated Server Time accuracy that is provided by STP is 100 milliseconds to the time provided by the NTP server. To meet the clock synchronization requirements of FINRA and MIFID II, the NTP server must have a pulse per second (PPS) output signal that can achieve time accuracy within 10 microseconds. If your configuration requires the NTP server with pulse per second capability, the NTP server that is configured as the ETS must be attached directly to the SE network.

For more information, see the Techdoc, *STP and FINRA clock synchronization requirements*.

8.4 Configuring the HMC as an NTP server

This section describes how to set up the HMC as an NTP server to be used as ETS for your CTN. The HMC can synchronize its time to an NTP server that is connected to the corporate network or available from the NTP pool on the internet.

The NTP server capability on the HMC addresses the potential security concerns that might arise if an external or internet NTP server is connected directly to the HMC/SE network. However, when you use the NTP server on the HMC as ETS for your CTN, no pulse-per-second capability is available.

In HMC Version 2.14.0, NTP Broadband Authentication that uses symmetric key (NTP V3-V4) and autokey (NTP V4) are supported. For more information about NTP Broadband Authentication, see 8.4.2, “NTP Broadband Authentication (optional)” on page 167.

8.4.1 Configuring HMC time source and NTP server

To configure the HMC time source to act as an NTP server, complete the following steps:

1. Open **Customize Console Date and Time** on the HMC. This task is available in the HMC Management section (see Figure 8-3).

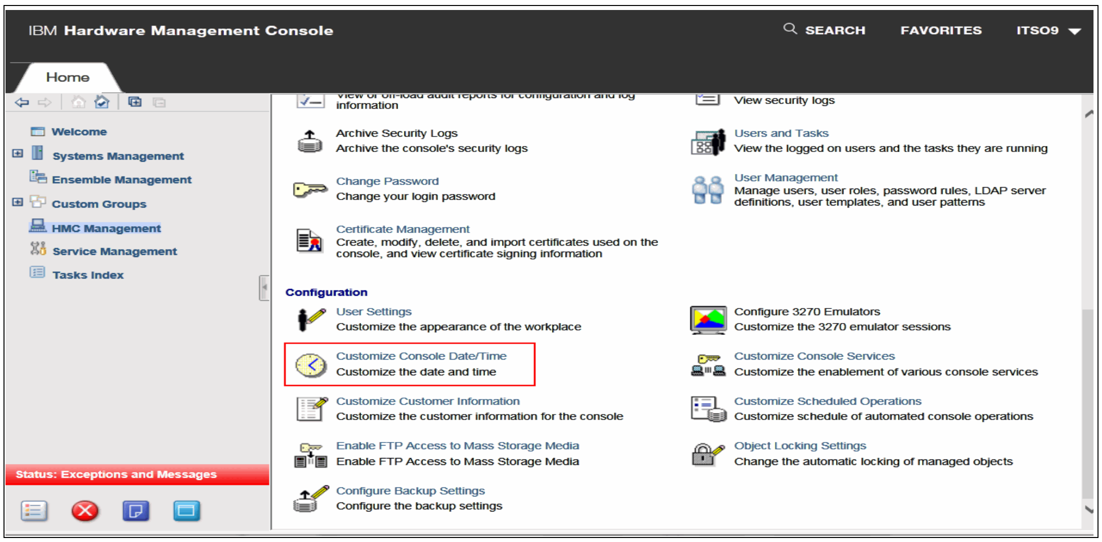

Figure 8-3 Customize Console Date and Time selection

2. The Date and Time window opens. Select **Network Time Protocol (NTP)** from Time Source (see Figure 8-4).

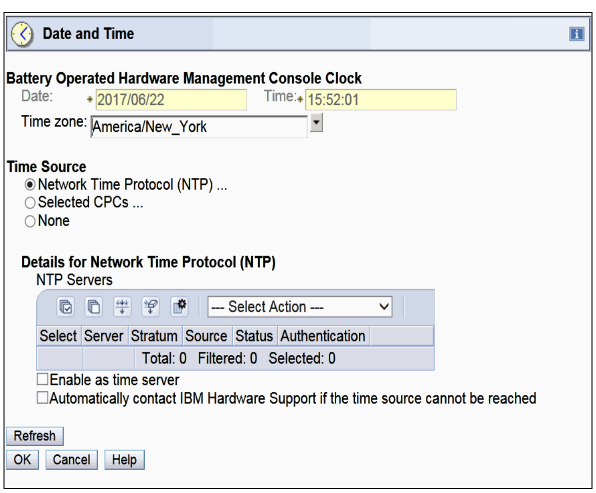

Figure 8-4 Customize Console Date and Time initial window

3. Select **Add Server** from the **Select Action** list (see Figure 8-5).

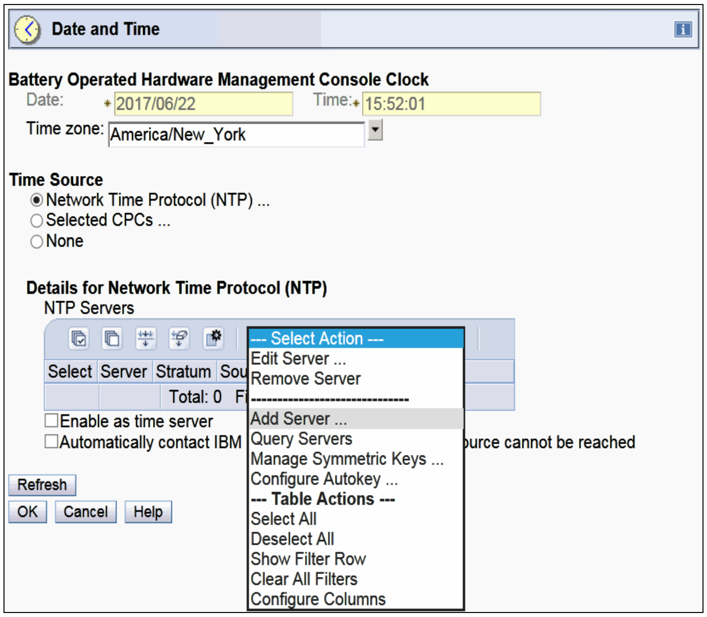

Figure 8-5 Selecting the Add Server option

4. The Add Network Time Server window opens. In this window, you can set the values in the following fields:
   - Enter the time server host name or IP address

     This mandatory field is the IP address or host name of NTP server. In our example configuration, `1.pool.ntp.org` is the NTP server.
   - Authentication Selection

     You can define this field if you want to use the NTP Broadband Authentication Support function. If you do not use the NTP Broadband Authentication Support function, select **None**.
5. Click **OK** when you are finished making your selections (see Figure 8-6).

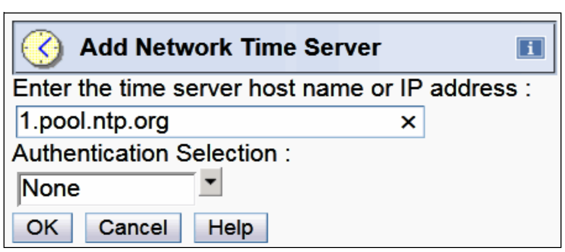

Figure 8-6 Add Network Time Server window

Tip: The use of a host name requires you to customize and enable Domain Name Services on the HMC. To complete this task, click **Customize Network Setting** → **Name services** in the Hardware Management Console Setting Work Area.

6. During this process, the HMC starts communicating with the NTP server that is specified. If NTP communication successfully completes, the window that is shown in Figure 8-7 is displayed. You can add multiple NTP servers by repeating step 4. Generally, define at least two NTP servers.

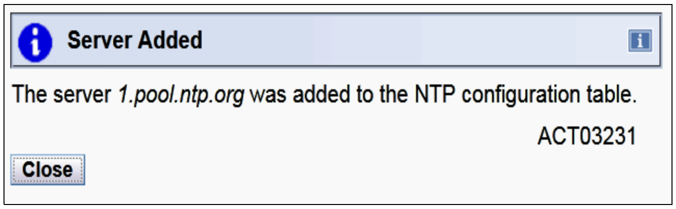

Figure 8-7 NTP server successfully added ACT03231 message

If the NTP server specified is not reachable, the Unable To Access Server window is displayed, as shown in Figure 8-8. Click **Yes** if you still want to add the NTP server.

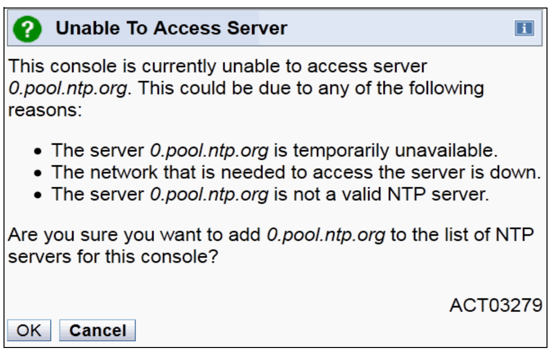

Figure 8-8 Unable To Access Server ACT03279 message

7. After NTP Servers are successfully added to the HMC, the Customize Console Date and Time window looks similar to the window that is shown in Figure 8-9. Select **Enable as time server** and click **OK** (see Figure 8-9).

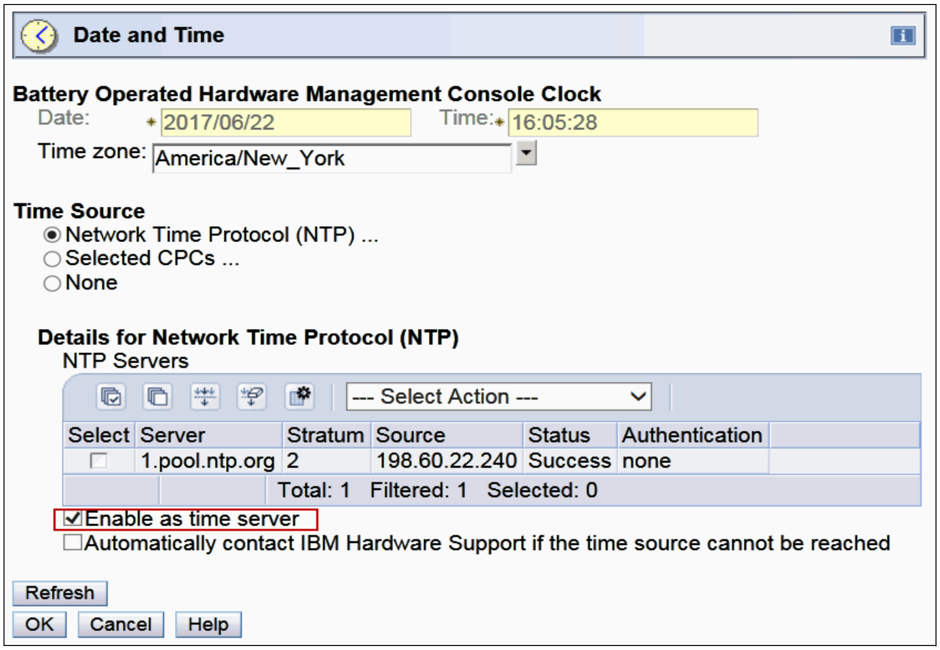

Figure 8-9 Configuring as NTP Server

8. In the Turning On NTP window (see Figure 8-10), confirm that you want to enable the HMC as time server. Click **OK**.

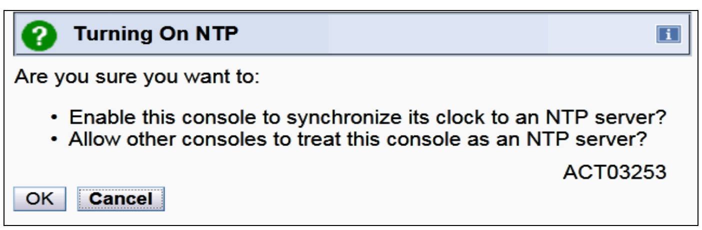

Figure 8-10 Turning On NTP server

9. After processing is completed, the NTP Successfully Turned On window opens (see Figure 8-11). Click **Close**.

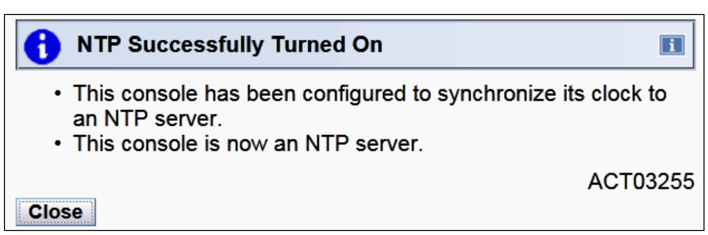

Figure 8-11 NTP Successfully Turned On window

Consideration: If more than one NTP server is defined, you cannot specify which server is the primary server. The NTP service on the HMC takes any defined NTP server and attempts to contact it. If it succeeds, that server is used as the time source until the server in question is no longer available or the console is restarted. If it cannot communicate with that server, it tries another server in the list.

The check box in the Configure NTP Setting window is used only for actions in the Select Action list, and *not* for setting a primary or preferred NTP server.

8.4.2 NTP Broadband Authentication (optional)

HMC supports NTP Broadband Authentication. Configure the (optional) NTP Authentication if your HMC requires it. NTP server authentication provides an increased level of security in the following situations:

- When a proxy is used to access an outside network

  Because NTP requests are User Datagram Protocol (UDP) socket packets, they cannot pass through the proxy. The proxy must be configured as an NTP server to get to target servers on the internet.

- When a firewall is used

  NTP requests can pass through the firewall. If you use a firewall to access an outside NTP server, use the HMC authentication to ensure untampered time stamps.

The following authentication methods are supported for NTP:

- Symmetric key (NTP V3-V4) authentication

  Symmetric key encryption uses the same key for encryption and decryption. When the HMC is acting as the client, the symmetric key index that is specified on each NTP server definition must be present in the key file.

  The specified key index, key type, and the key string must align with the specified key information of the target server. Likewise, if the HMC is acting as a server, the client specified key information must match the same key index on the server. Symmetric key supports Network Address Translation (NAT).

- Autokey (NTP V4) authentication

  Autokey uses public key cryptography. The key generation for the HMC NTP is done by clicking **Generate Local Host Key** in the Autokey Configuration window. Clicking this option runs the `ntp-keygen` command, which generates the specific key and certificate for this system. Autokey authentication is not available with a NAT firewall.

Setting up a Symmetric key

To set up a Symmetric key, complete the following steps:

1. Open the Customize Console Date/Time window, which is described in step 1 on page 163.
2. Click **Select Action** → **Manage Symmetric keys** in the Details for Network Time Protocol (NTP) section, as shown in Figure 8-12.

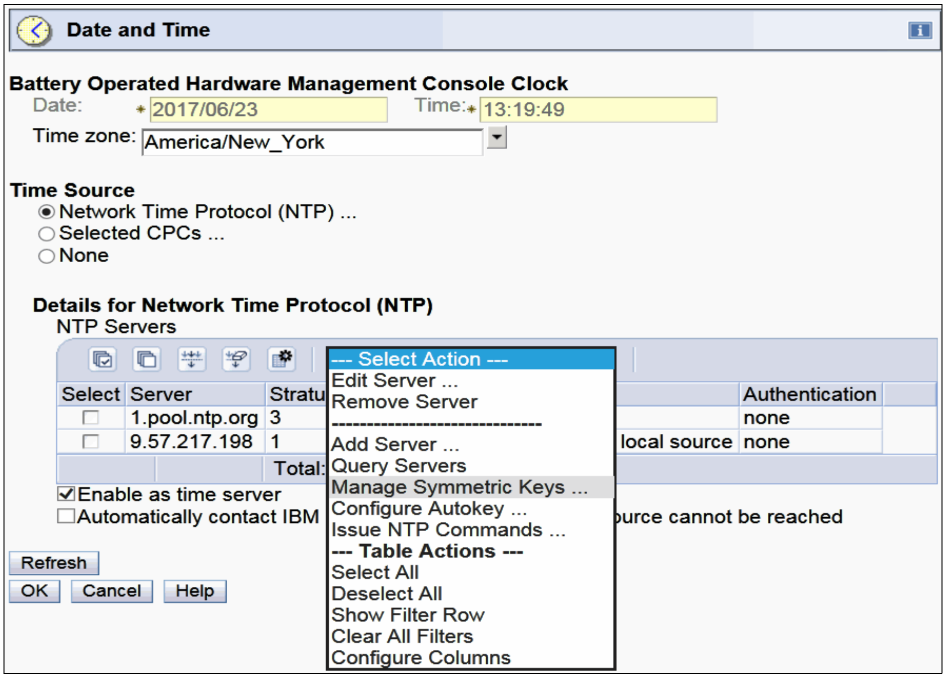

Figure 8-12 Selecting the Manage Symmetric Keys option

3. The Manage Symmetric Keys window opens (see Figure 8-13). Click **Select Action** → **Add Key**.

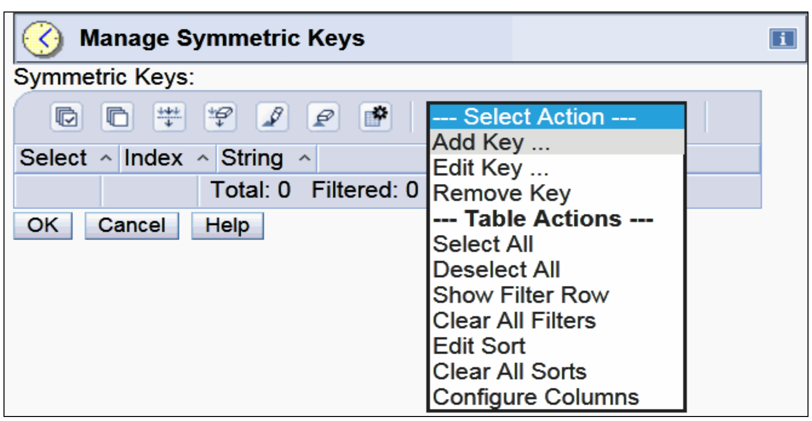

Figure 8-13 Selecting the Add Symmetric Key option

4. The Add Symmetric Key Data window opens (see Figure 8-14). The following data must be entered:

Key index A numeric value in the range 1 - 65534.

Key string The key string can be up to 40 characters long. If the string is 40 characters long, the characters must be hexadecimal ASCII characters (0 - 9, a - f). If the string is fewer than 40 characters long, the characters can be any printable ASCII character.

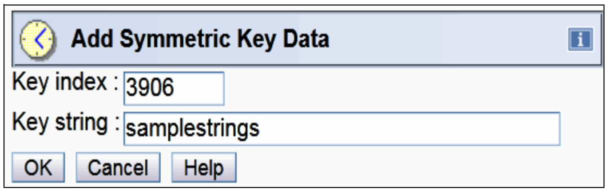

Figure 8-14 Add Symmetric Key Data

5. The Enter Symmetric Key window opens. Click **Close** to add the Symmetric Key (see Figure 8-15).

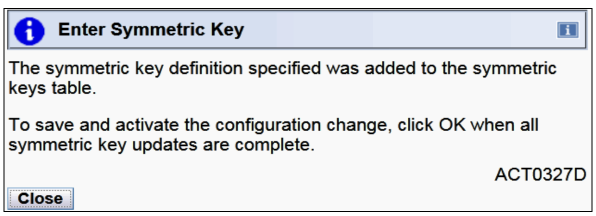

Figure 8-15 Enter Symmetric Key successfully added

6. The key is displayed in the Manage Symmetric Keys window (see Figure 8-16). You can enter more key data in this window by repeating the previous steps. Click **OK** to proceed if no other keys are needed.

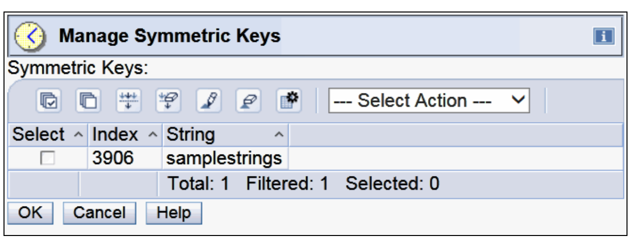

Figure 8-16 Symmetric Key list

7. The completion window opens (see Figure 8-17). Click **Close** to return to the Configure NTP Settings window.

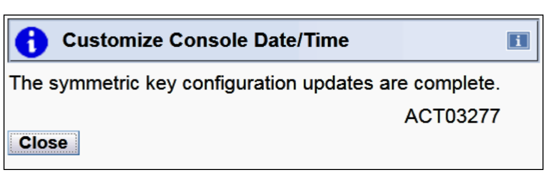

Figure 8-17 Completed message for symmetric key configuration

8. You can define a Symmetric Key to NTP server by using one of the following methods:
   – Define a new NTP server and specify a Symmetric Key
   – Modify an existing NTP server to use a Symmetric Key

   For this example, modify the NTP server that is identified by 9.12.5.155 (IP address) with a Symmetric Key.

9. Select the server name that you want to modify and then, click **Select Action** → **Edit Server** (see Figure 8-18).

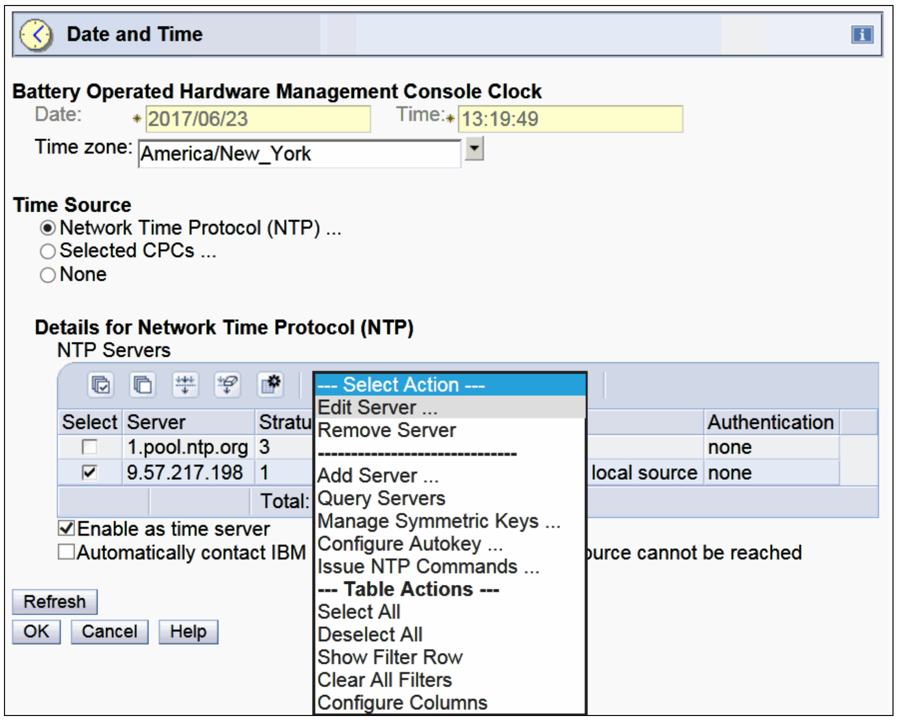

Figure 8-18 Selecting the Edit Server option

10. In the Edit Network Time Server window (see Figure 8-19), select the appropriate key number in the **Symmetric Key** field. In this example, we select key 3906. Click **OK** to continue.

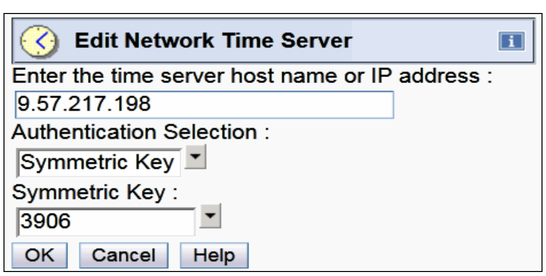

Figure 8-19 Adding NTP Server with Authentication

11. The next window (see Figure 8-20) indicates that the modification is completed. Click **Close**.

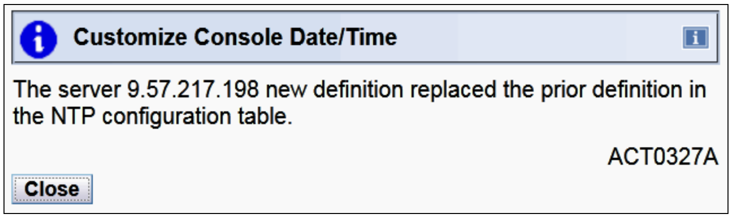

Figure 8-20 Completed message for adding NTP Server with Authentication

12. You are returned to the Date and Time window (see Figure 8-12 on page 168). Ensure that the key that you selected is displayed in the Authentication column, the Status is `Success`, and the key that you specified is displayed.

Configuring NTP authentication using Autokey

To configure Autokey, complete the following steps:

1. Open the **Customize Console Date/Time** window, which is described in step 1 on page 163.
2. Under the Details for Network Time Protocol (NTP) section, click **Select Action** → **Configure Autokey** (see Figure 8-21).

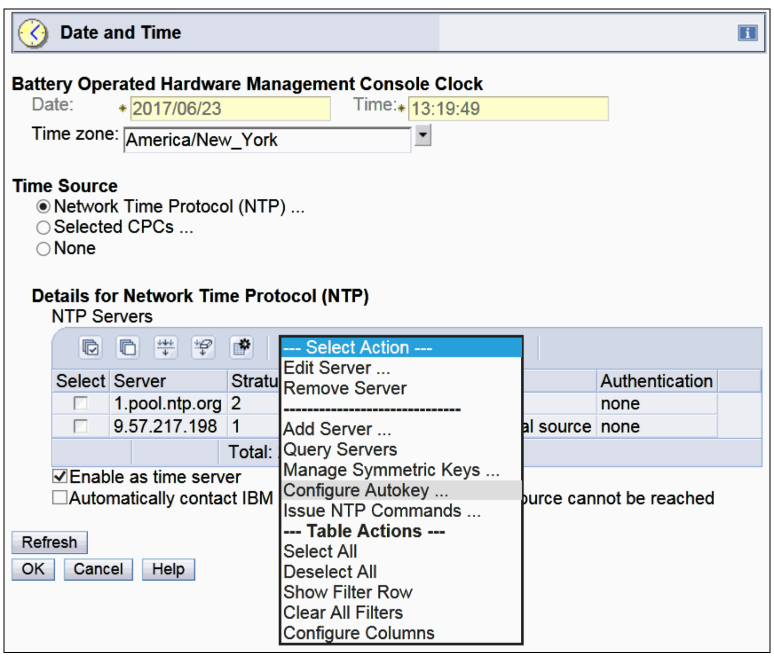

Figure 8-21 Selecting the Configure Autokey option

3. The Autokey Configuration window opens (see Figure 8-22). To generate keys, click **Generate**.

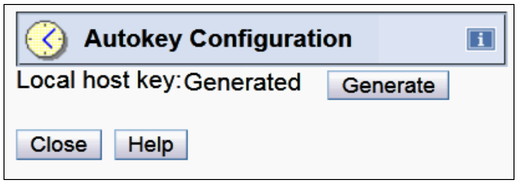

Figure 8-22 Generate Autokey

4. When the Autokey is successfully generated, the confirmation window (see Figure 8-23) opens. Click **Close**.

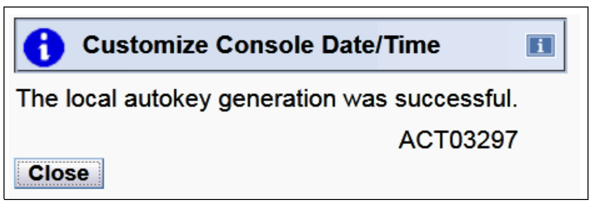

Figure 8-23 Autokey generation successful

5. You can define an Autokey to NTP in the following cases:
   - Defining a new NTP server and specify an Autokey.
   - Modifying an existing NTP server to use an Autokey.

   In this example (see Figure 8-24), we modify the existing NTP server `9.57.217.198` to use Autokey configuration.

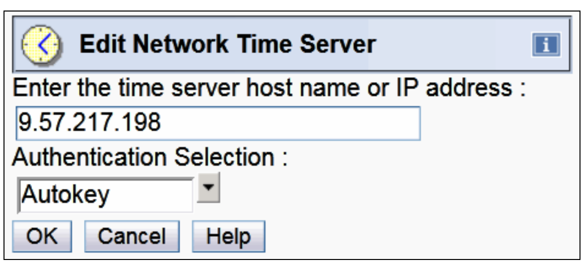

Figure 8-24 Edit Network Time Server for Autokey

6. Select the server that you want to modify, then click **Select Action** → **Edit Server**.
7. The Edit Network Time Server window opens (see Figure 8-24). Select **Autokey** under Authentication Selection. Then, click **OK**.
8. The next window (see Figure 8-25) indicates that the modification completed successfully. Click **Close**.

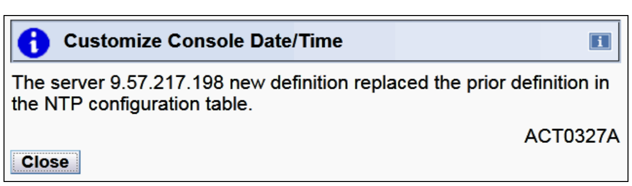

Figure 8-25 NTP setting change confirmation window

9. Return to the Configure NTP Setting window. Ensure that `autokey` is listed in the Authentication column, and its status is `Success`.

8.5 New in HMC V2.14.0: Manage System Time task

You can set up STP by using the HMC **Manage System Time** task. HMC Version 2.14.0 brings the following changes:

- HMC task menu title is changed from System(Sysplex) time to Manage System Time.
- CTN topology is displayed as a graph.
- Wizard format operations are supplied for STP actions, such as setting up the CTN and changing the STP server roles.

Figure 8-26 shows the **Manage System Time** task initial window on the HMC. 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 that is next to the CTN ID name. You can enter the configuration Wizard from the menu that is below **STP ACTIONS**.

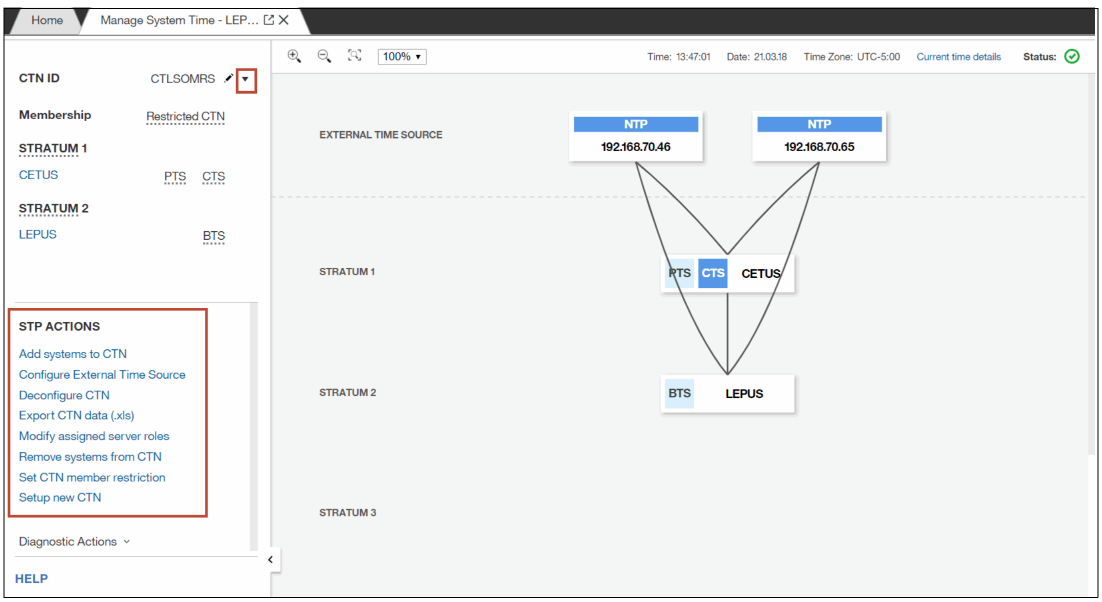

Figure 8-26 Manage System Time initial window

In the Manage System Time initial window, the CTN topology displays graphically. 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.

Figure 8-27 shows the details of a CPC in the CTN.

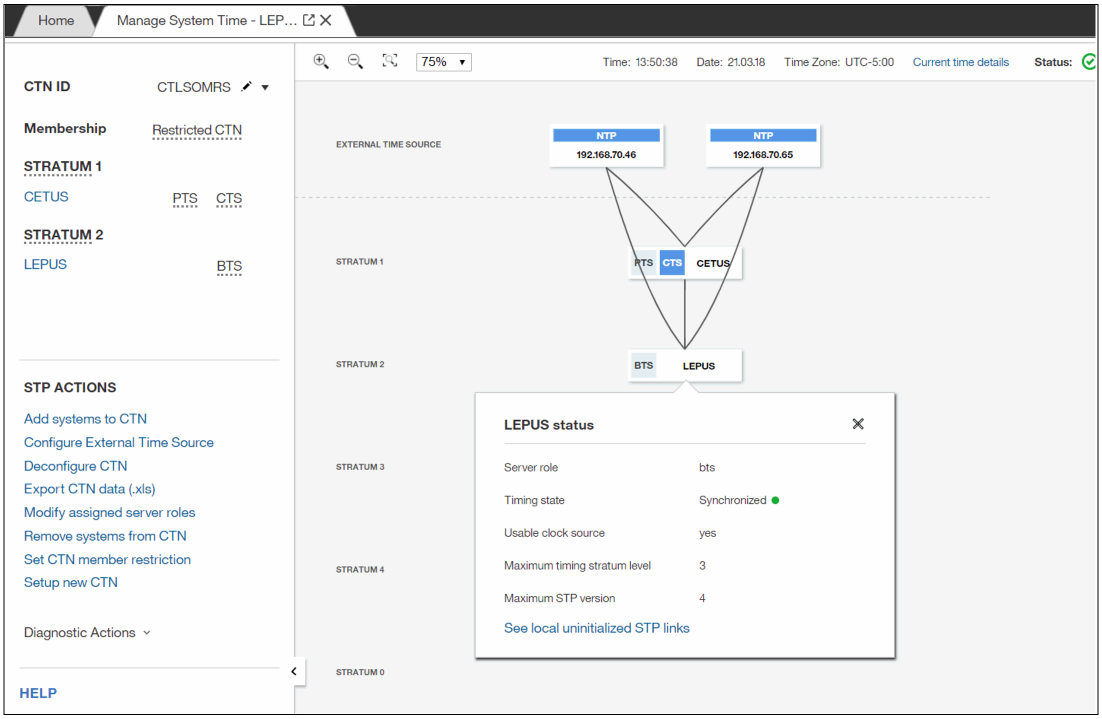

Figure 8-27 CPC status display

Figure 8-28 shows the details of a coupling link.

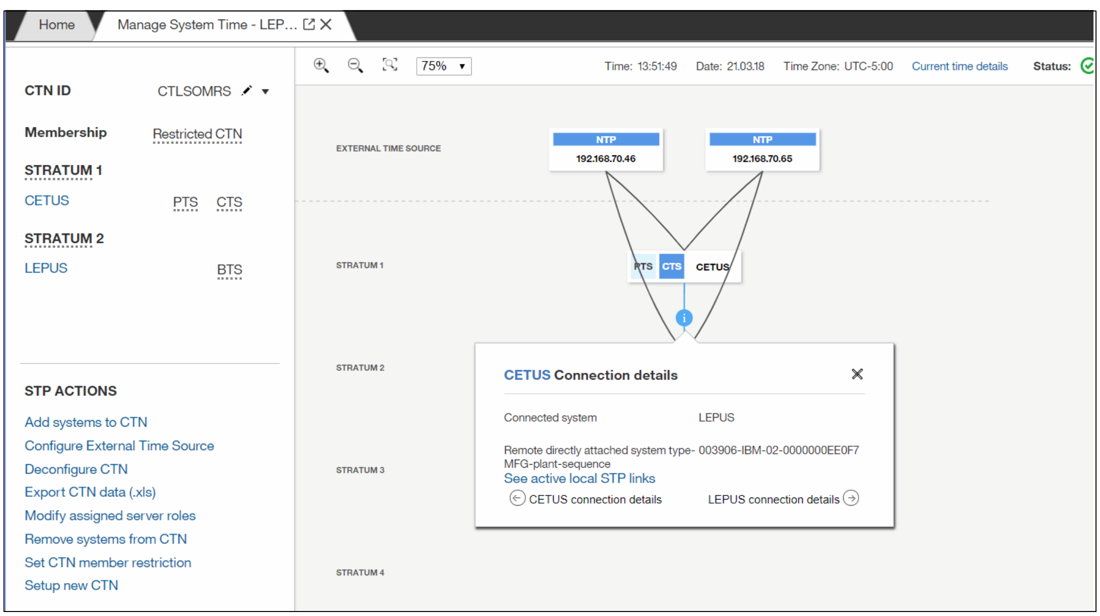

Figure 8-28 Coupling/timing link status display

Note: In z14 SE, the System (Sysplex) Time menu still exists[a]. The task still uses same interface that was provided in previous HMC (2.13.1). If you use SE for setting up STP, see *IBM z13® Configuration Setup*, SG24-8260.

a. For consistency, use the HMC interface for STP-related tasks.

8.5.1 Setting CTN member restriction menu

From HMC V2.14.0, a new menu, Set CTN member restriction, was added. You can set the Only allow the server(s) specified above to be in the CTN parameter 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).

To set the Only allow the server(s) specified above to be in the CTN parameter, complete the following steps:

1. Select **Set CTN member restriction** from STP ACTIONS (see Figure 8-29).

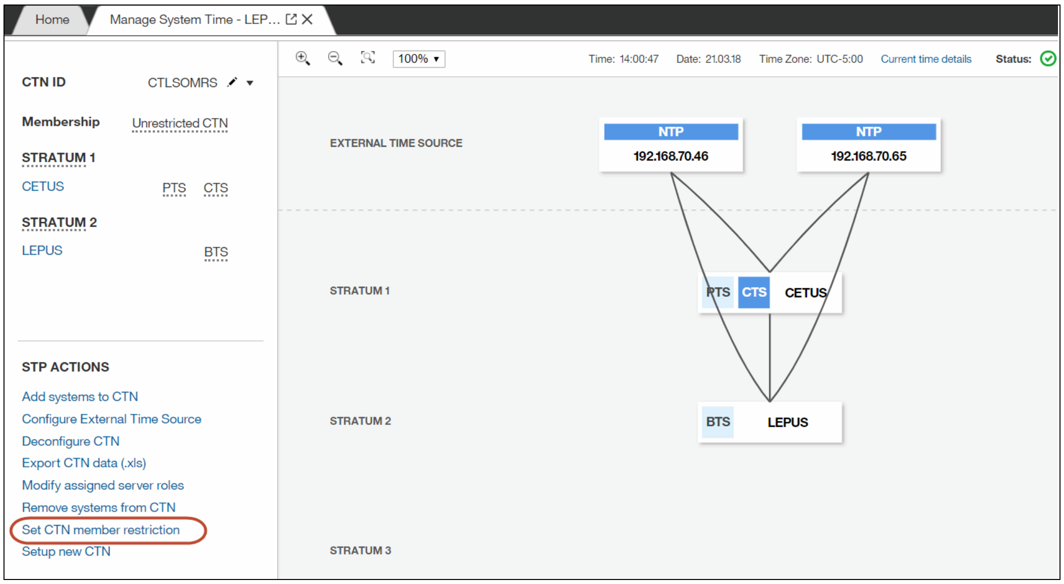

Figure 8-29 Set CTN member restriction

2. Select **Allow only the servers that are specified below to be members of the CTN** and click **APPLY** (see Figure 8-30). Confirm that the Success applying changes message is displayed after APPLY is clicked.

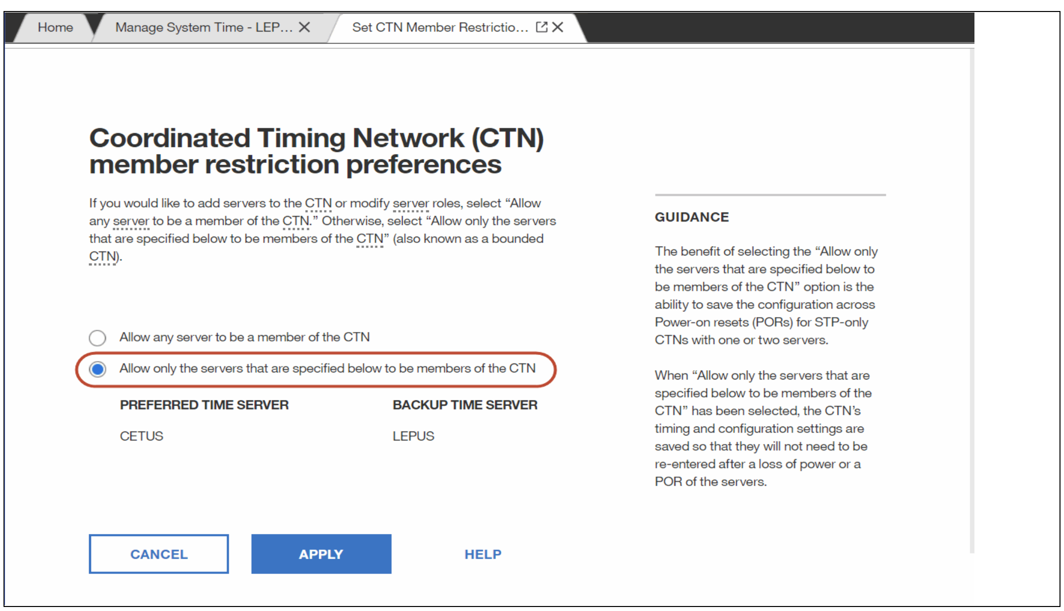

Figure 8-30 CTN member restriction preferences

8.6 Single-server STP-only CTN

This section describes how to configure new STP-only CTN. This scenario applies to a newly installed z14 or isolated STP-only CTN for testing purposes.

This example defines a STP-only CTN named ITSO. In this CTN configuration, we define z14 (CETUS) and two NTP servers as ETS. The NTP servers to be used as ETS are defined to the HMC (see Figure 8-31). This configuration is a single CPC STP-only CTN, so no BTS and ARBITER are configured.

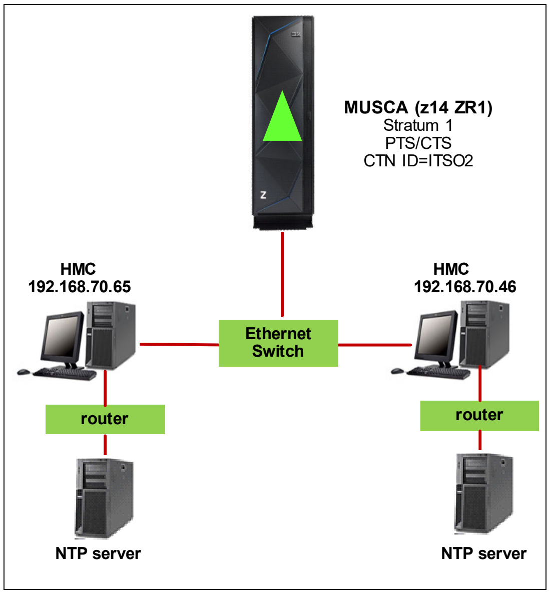

Figure 8-31 STP-only CTN configuration, single CPC example

Note: Before you configure STP, see the following publications:

- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Implementation Guide*, SG24-7281
- *Server Time Protocol Recovery Guide*, SG24-7380

8.6.1 Configuring a new STP-only CTN

To configure a new STP-only CTN, select the **Setup new CTN** menu from the STP ACTIONS list in the **Manage System Time** task. To set up new STP-only CTN, complete the following steps:

1. In the HMC, select the CPC to configure STP and **Manage System Time** under the Configuration task (see Figure 8-32).

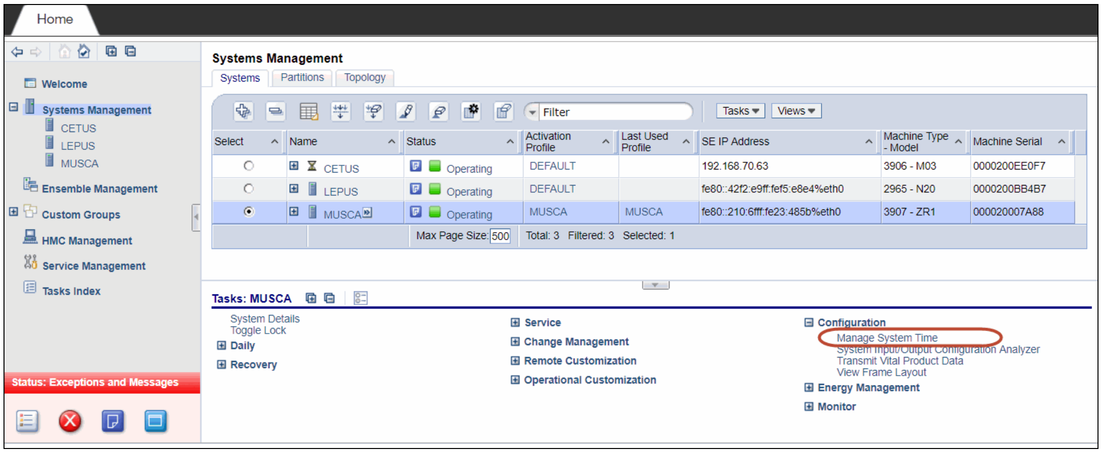

Figure 8-32 Manage System Time CPC selection

2. The Manage system Time window opens. If defined CPC objects exist that configured a CTN, the CTN configuration is displayed graphically. To define a new CTN, select **Setup new CTN** from the STP ACTIONS list (see Figure 8-33).

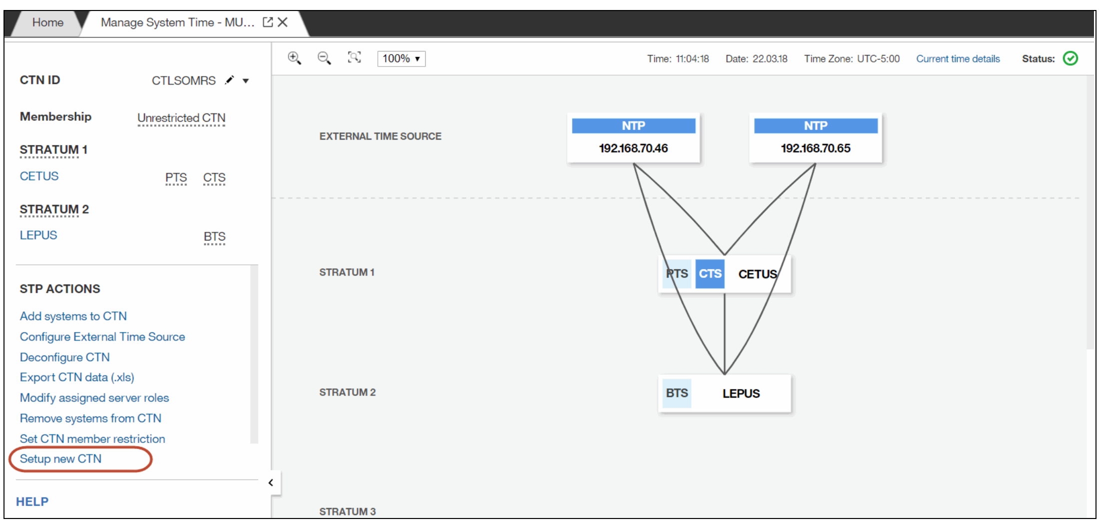

Figure 8-33 Setting up new CTN

3. The Set the Coordinated Timing Network (CTN) ID window opens. Enter the CTN ID in the column. Click **NEXT** to continue (see Figure 8-34).

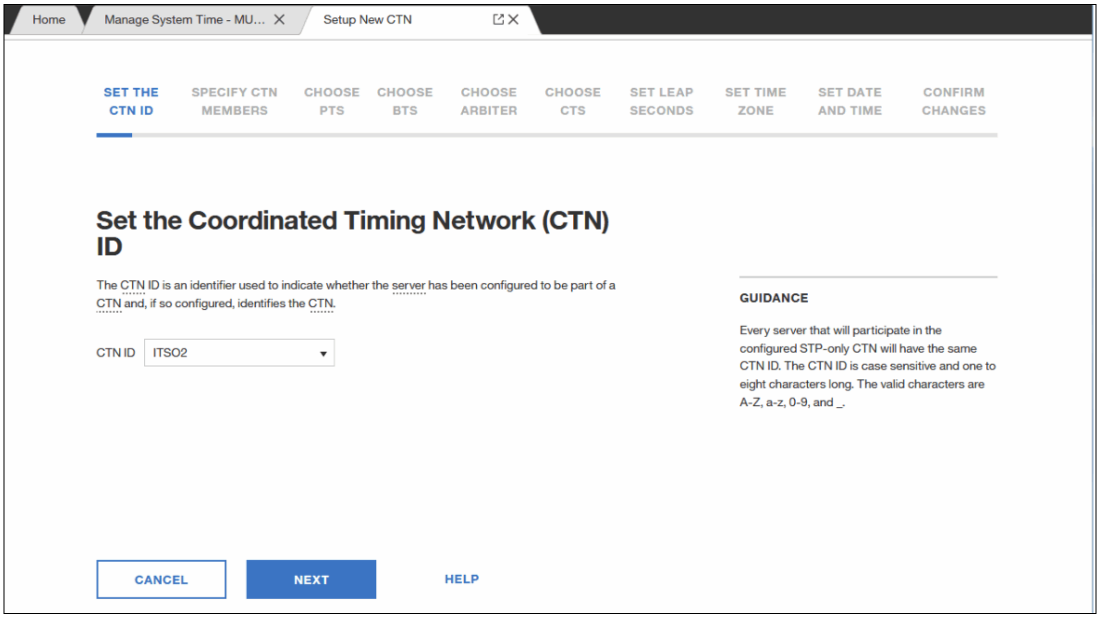

Figure 8-34 Setting the Coordinated Timing Network (CTN) ID

4. The Specify Coordinated Timing Network (CTN) members window opens. Select the CPC name (only CPCs that do not belong to a CTN are shown). In this example, only one CPC (CETUS) is displayed. Select the CPC and click **NEXT** to continue (see Figure 8-35).

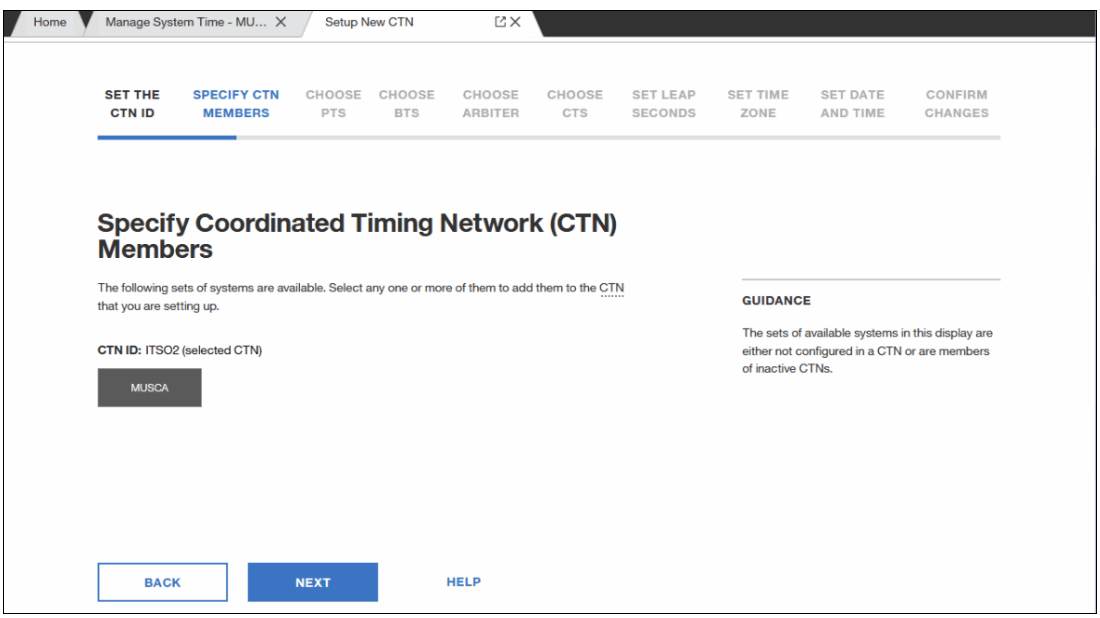

Figure 8-35 Specify Coordinated Timing Network (CTN) members

5. The Choose Preferred Time Server window opens. Select the CPC name for PTS (see Figure 8-36).

> **Note:** Because MUSCA is the only CPC (server) in the CTN, the CHOOSE BTS and CHOOSE ARBITER steps are not available.

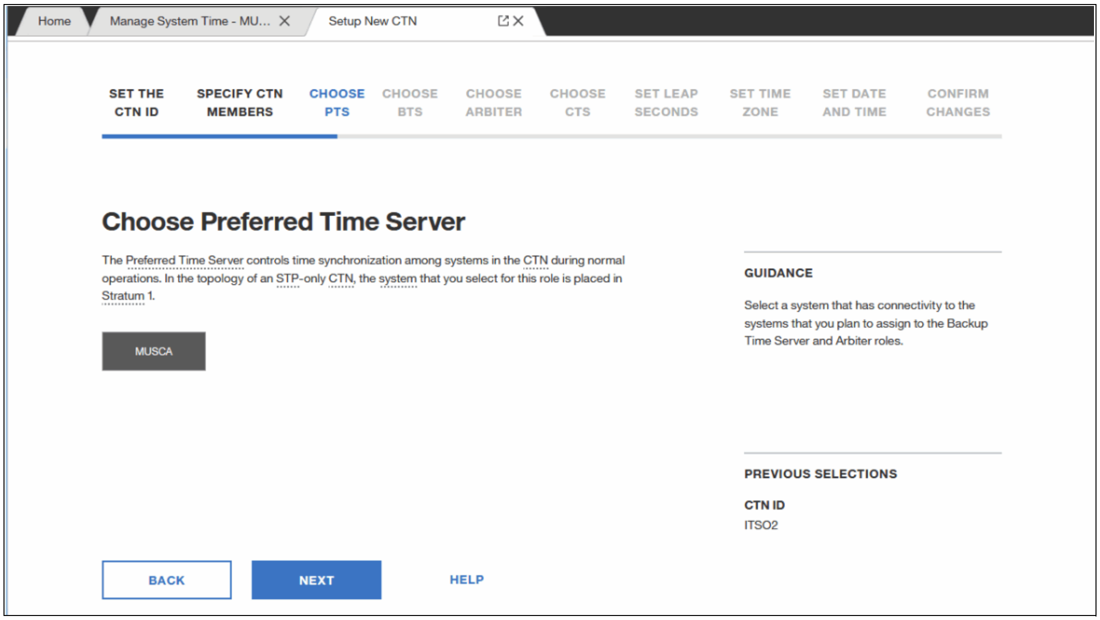

Figure 8-36 Choose Preferred Time Server window

6. The Choose Current Time Server window opens. We select **MUSCA** for CTS. Click **NEXT** to continue (see Figure 8-37).

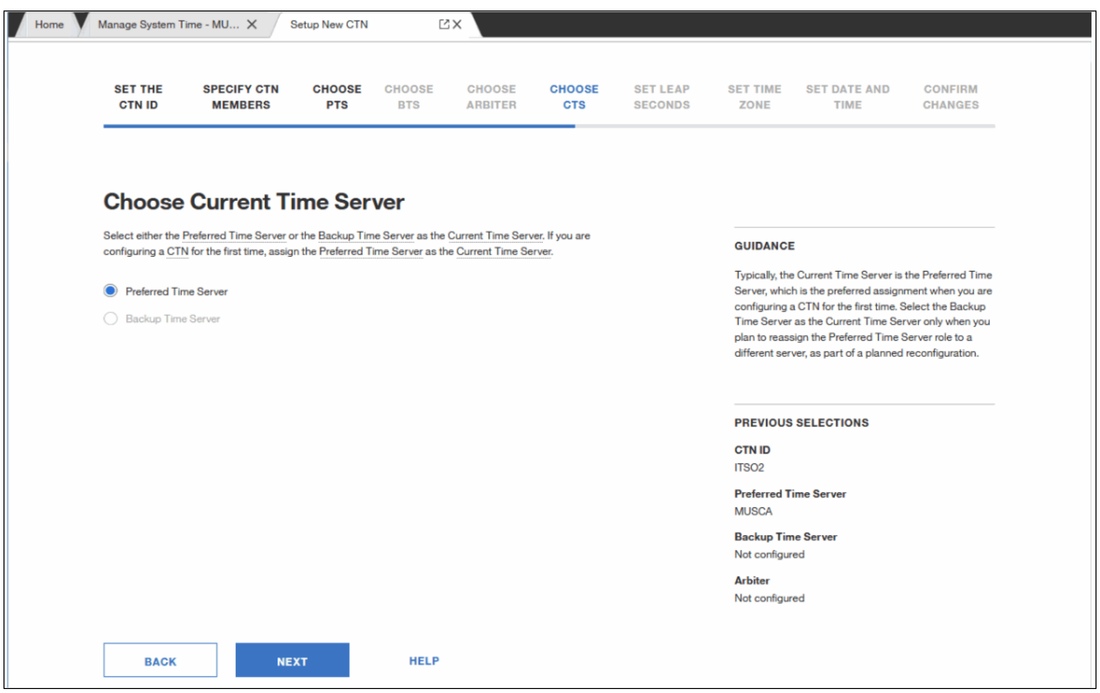

Figure 8-37 Choose Current Time Server window

7. The Set leap second window opens. Define the applicable value for the leap seconds offset. Then, click **NEXT** to continue (see Figure 8-38).

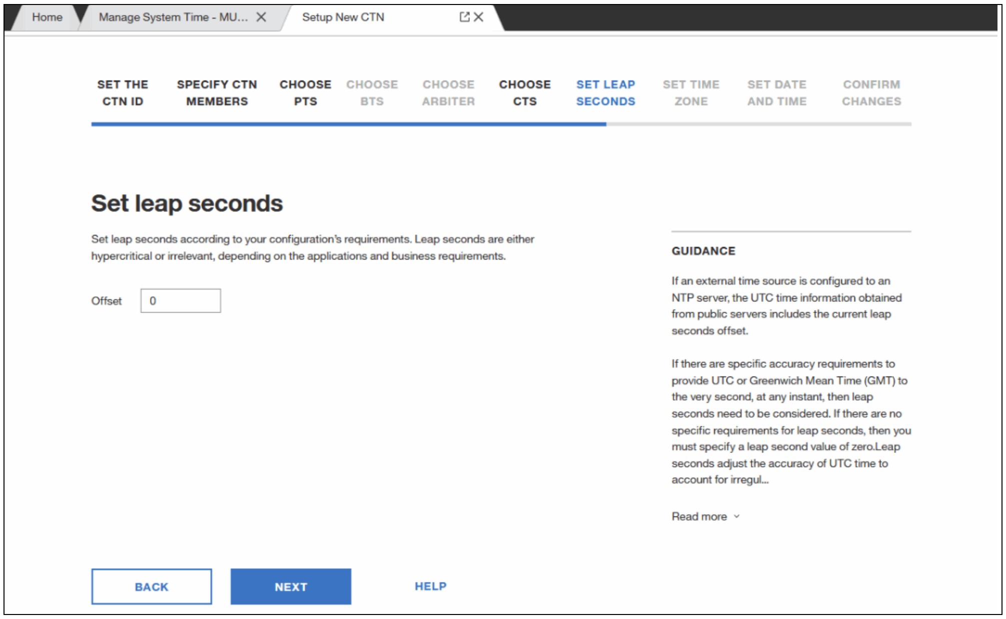

Figure 8-38 Set leap seconds

8. The Set time zone window opens. Specify the following parameters:

   – Adjust time zone offset
   – Clock adjustment for Daylight Saving Time

   Select the appropriate items from the list. These parameters are used when time is initialized for the CTS that is defined in the next step. We define **Eastern Time** for the Time Zone offset and **Automatically adjust** for the Daylight Saving Time. Click **NEXT** to continue (see Figure 8-39).

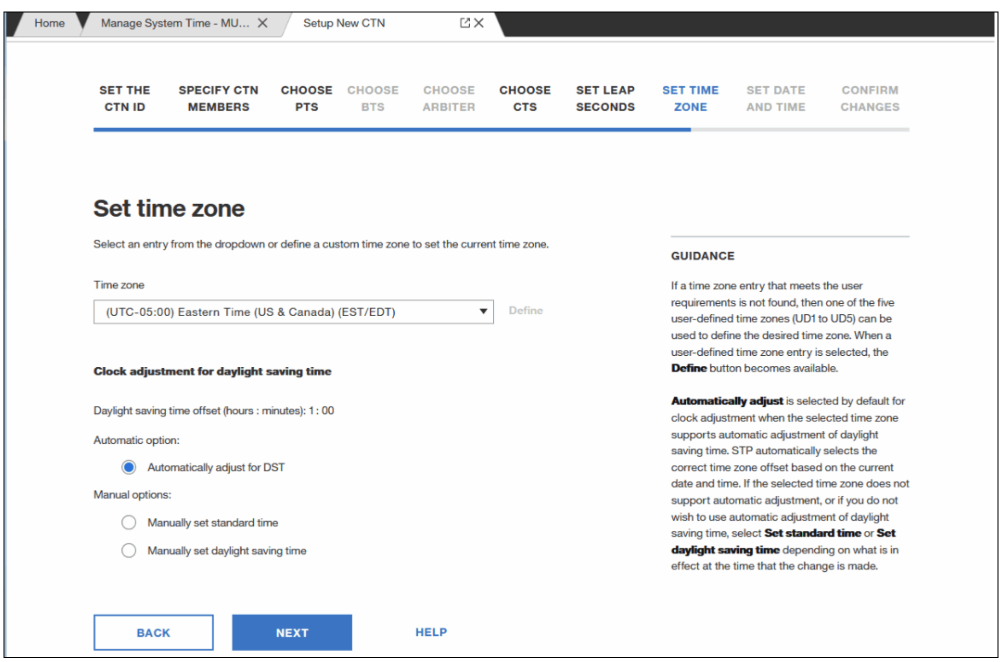

Figure 8-39 Set time zone

9. The Set date and time window for initializing the time opens. Select one of the following three items to initialize the time:
   - Use the configured External Time source to set date and time

     You can set ETS during the initializing time process.
   - Set date and time

     Select this option when you set the date and time for CTS manually.
   - Modify time by delta to set date and time

     Select this option when you set the current TOD of CTS for initializing the time and specify a delta from the TOD.

   We select **Use the configured External Time source to set date and time**, and click **NEXT** to continue (see Figure 8-40).

   > **Note:** Selecting **Use the configured External Time source to set date and time** enables accurate setting of the initial time; therefore, use this option in general.

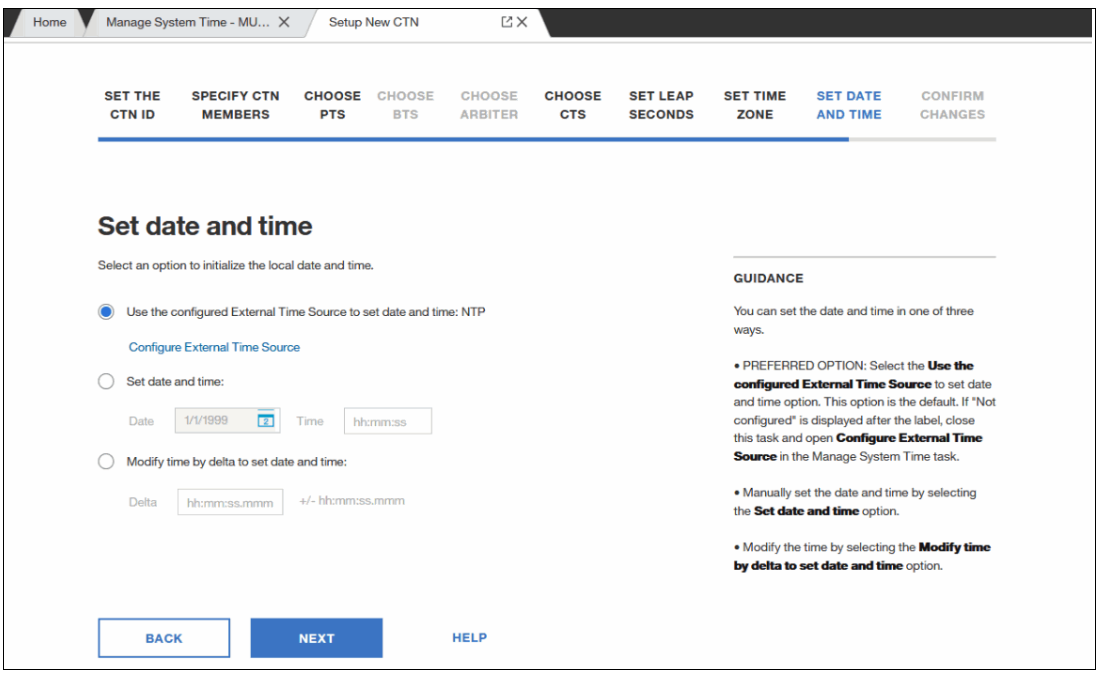

Figure 8-40 Set date and time

10.The Confirm Changes window opens (see Figure 8-41). The topology of CTN you defined is displayed graphically.

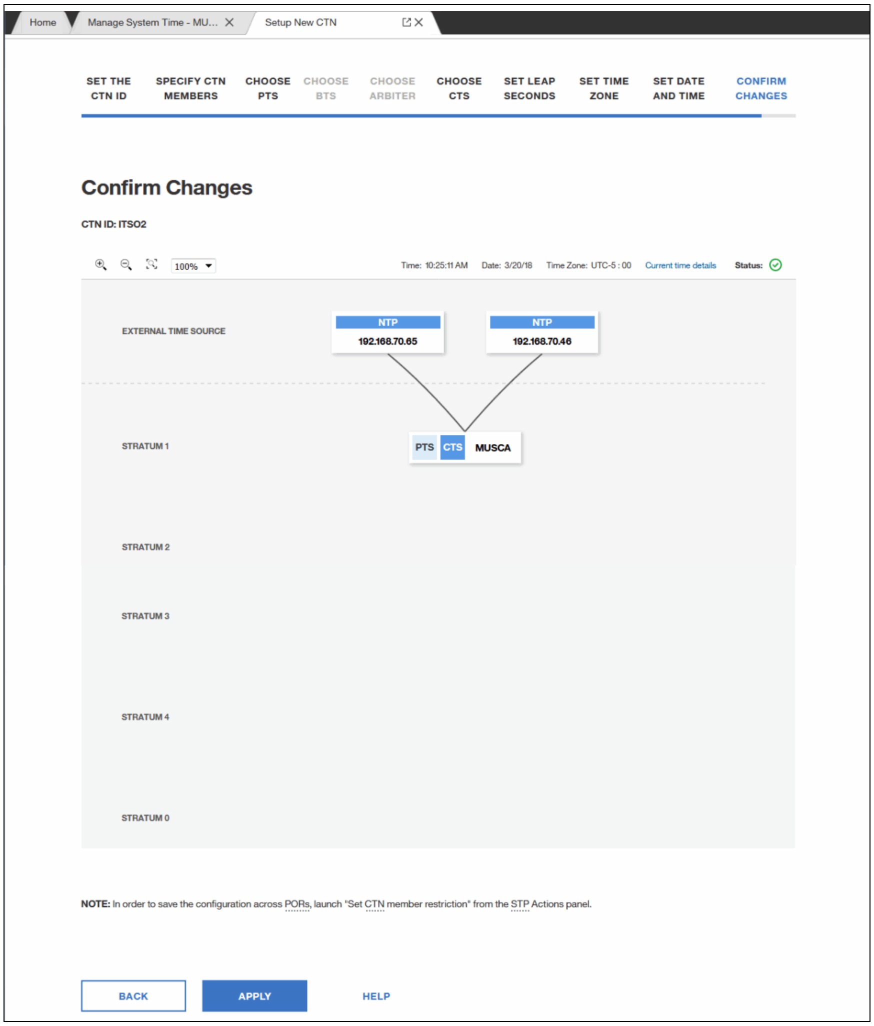

Figure 8-41 Confirm Changes window

11.The Creating CTN progress is shown until the operation completes successfully, followed by the ACT39277I success window (see Figure 8-42). Click **Close** to return to the initial window.

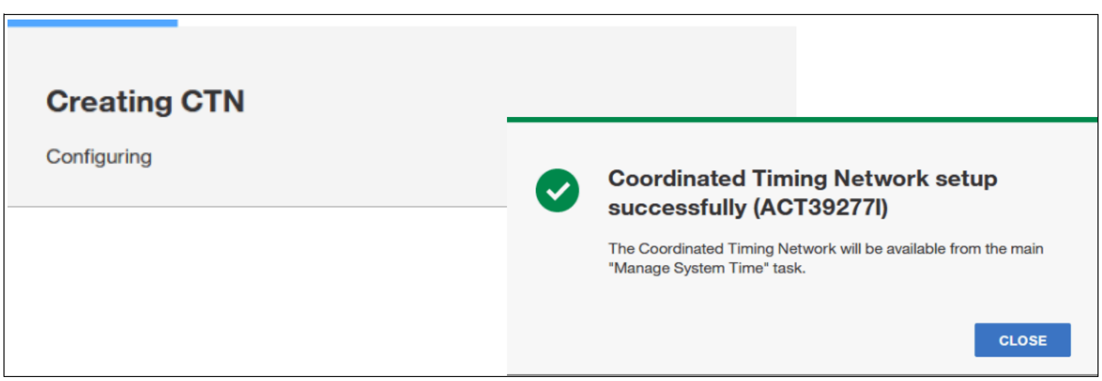

Figure 8-42 Creating CTN window / ACT39277I Success window

8.6.2 Verifying new CTN configuration

When the new CTN configuration is complete, the Manage System Time initial window is displayed. Check that the following items are defined properly on this window (see Figure 8-43):

- Status
- CTN ID
- Time/Date/Time zone
- Stratum Level
- STP Role

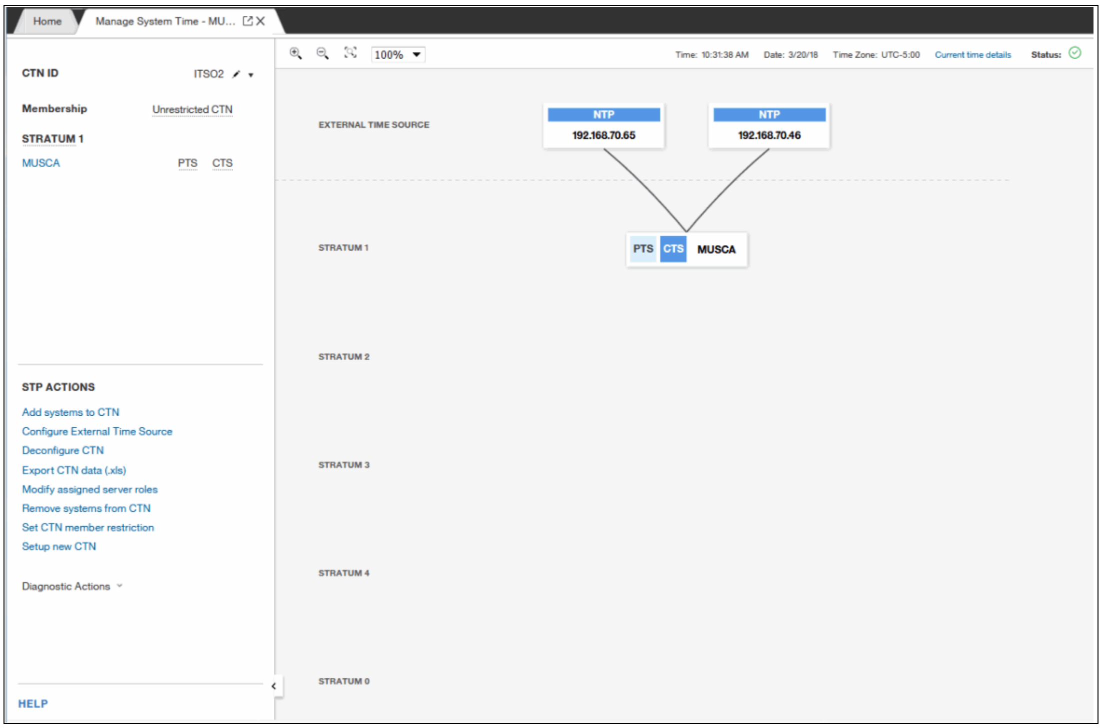

Figure 8-43 New CTN

You can also check the status of STP by using the z/OS **D ETR** command (see Figure 8-44).

```
D ETR
IEA386I 13.15.59 TIMING STATUS 951
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = IZTSO2
  THE STRATUM 1 NODE ID = 003907.ZR1.IBM.02.000000007A88
  THIS IS THE PREFERRED TIME SERVER
```

Figure 8-44 z/OS D ETR command

8.6.3 Configuring External Time Source

After the new CTN configuration process is completed, you can configure the ETS. In this example, we define two HMCs that are configured as NTP server for ETS (see Figure 8-33 on page 177). If your CTN is configured with more than two CPCs, you can set individual ETS for PTS and BTS.

Complete the following steps:

1. Open the **Manage System Time** task and select **Configure External Time Source** from the STP ACTIONS list.
2. The Select the system on which to modify the External Time Source (ETS) window opens. Select the CPC for which to modify the ETS setting. We select **MUSCA** (the only CPC in this CTN) in this window. Click **NEXT** to continue (see Figure 8-45).

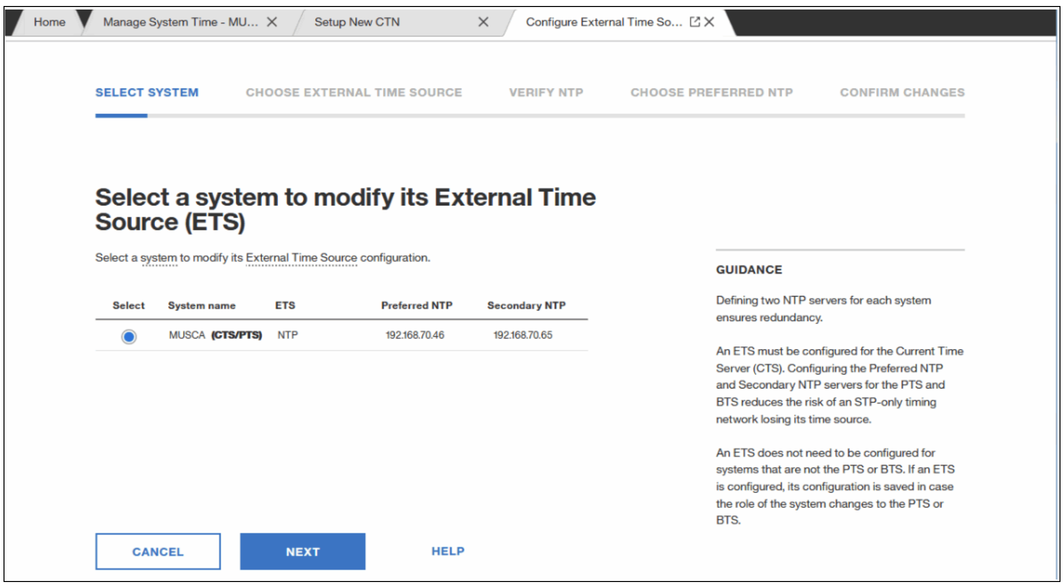

Figure 8-45 Select the system on which to modify the External Time Source (ETS) window

3. The Choose External Time Source window opens. We select **Use NTP**. Click **NEXT** to continue (see Figure 8-46).

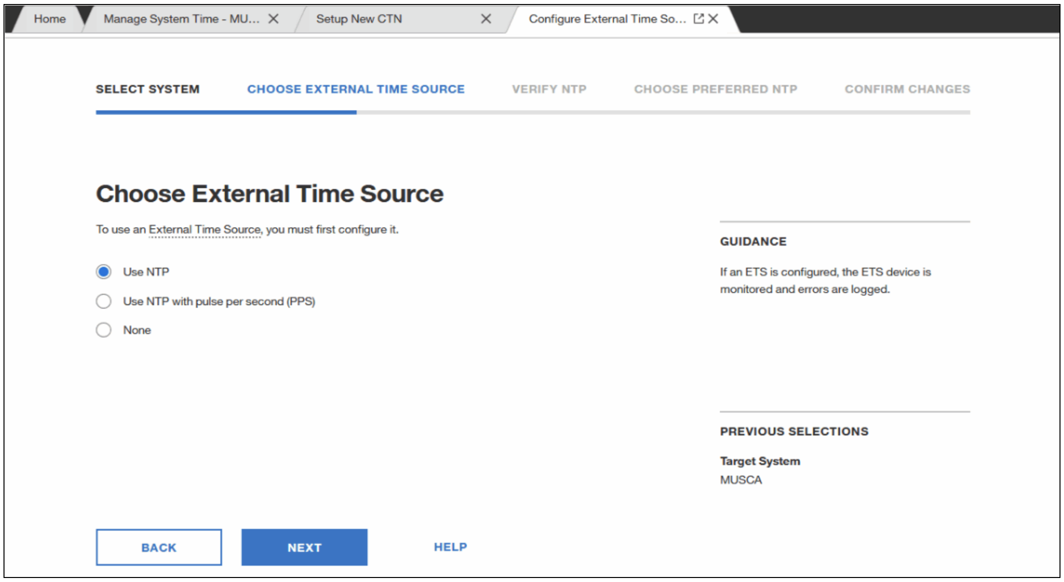

Figure 8-46 Choose External Time Source window

4. The Verify Network Time Protocol servers window opens (see Figure 8-47). You can define up to two NTP servers by specifying the IP address or the host name (IP label) of the NTP server by clicking the pencil icon. After IP address or host name is added, test the connectivity to ensure that the Enabled switch icon is enabled. Then, click **TEST CONNECTIVITY** to ensure that the NTP server is reachable. If the connectivity test is successful, the ACT3929I window opens and the Check the Connection status row displays `No errors`. Click **NEXT** to continue.

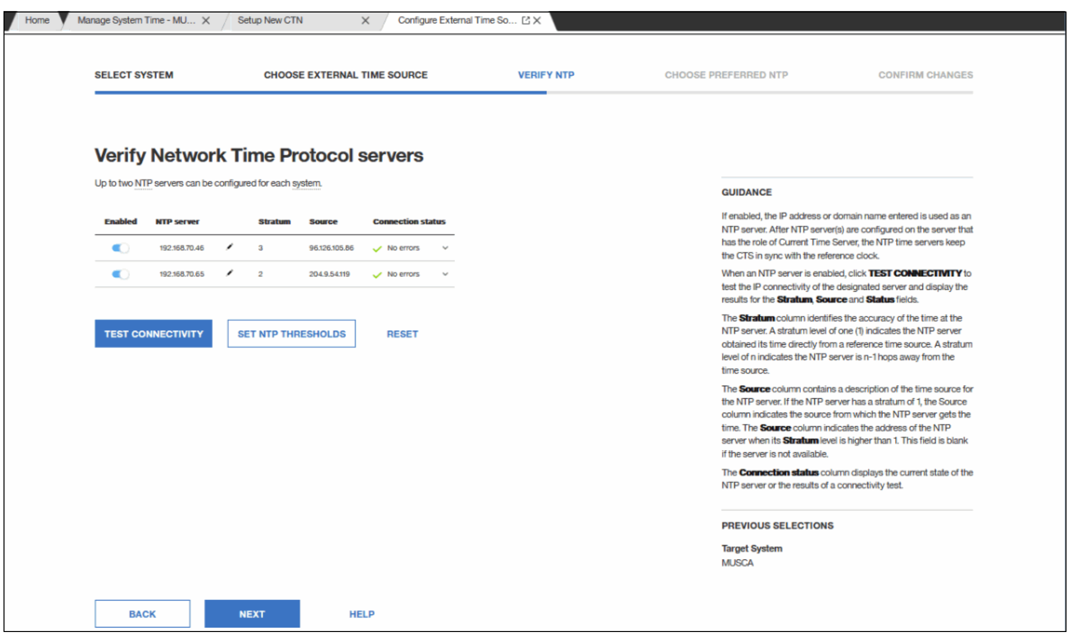

Figure 8-47 Verify Network Time Protocol servers

If you need to set NTP thresholds, click **Set NTP Thresholds** and specify the parameter. You can set the NTP Stratum level threshold and the Source ID time threshold (see Figure 8-48).

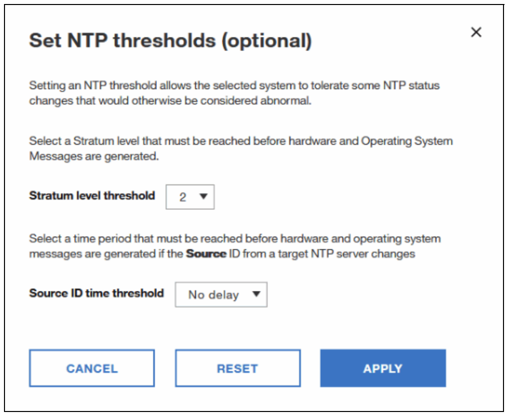

Figure 8-48 Set NTP Thresholds (optional)

5. The Choose the preferred NTP server window opens (see Figure 8-49). Select one of the NTP servers to be the preferred server. In our example, we select **192.168.70.46** as the preferred NTP server. Click **NEXT** to continue.

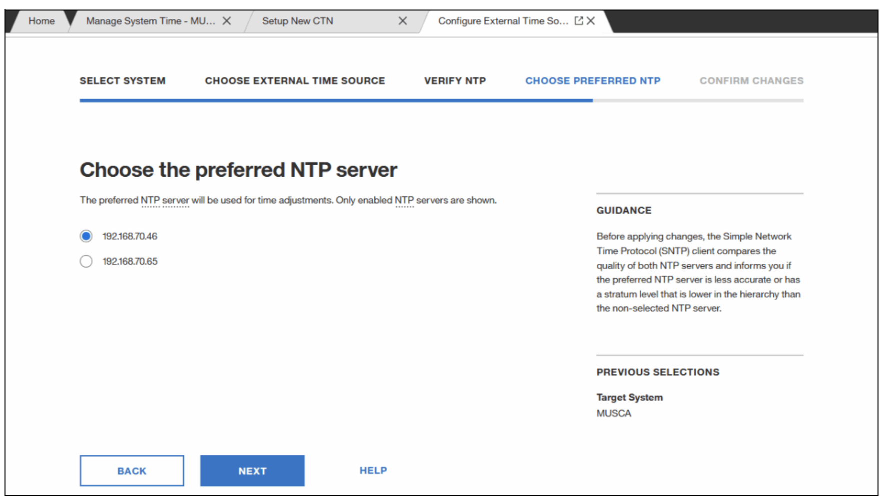

Figure 8-49 Choose the preferred NTP server window

6. The Confirm External Time Source configuration window opens. Verify and click **APPLY** to complete the definition (see Figure 8-50).

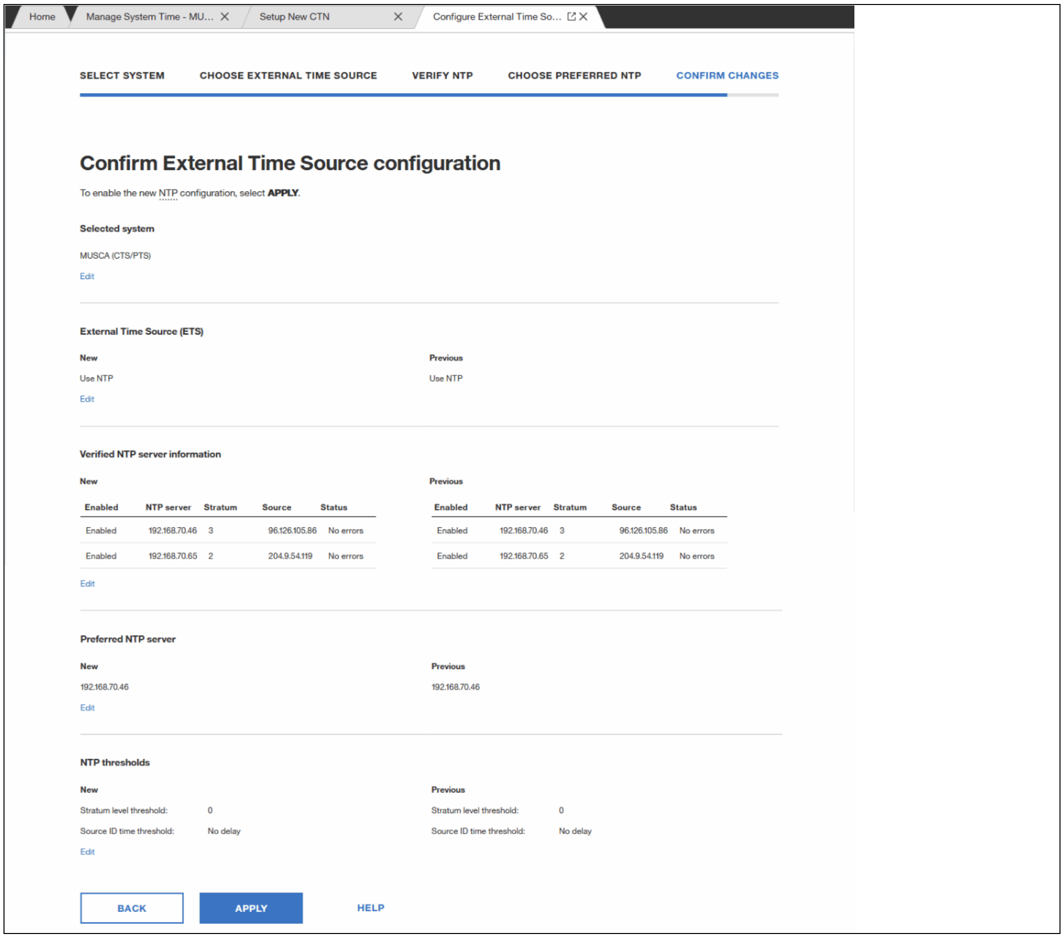

Figure 8-50 Confirm External Time Source configuration window

7. The Set ETS configuration (ACT 39145) window opens. Click **DONE** to finish. (Figure 8-51).

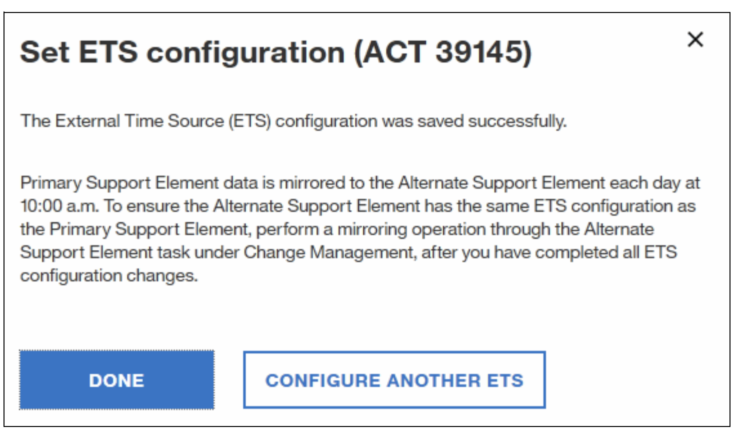

Figure 8-51 ETS configured successfully (ACT39145)

If no changes were made, the ACT39290I window opens in which this information is displayed.

8.6.4 Verifying the ETS configuration

From the Manage System Time initial window, check whether the defined ETS is displayed (EXTERNAL TIME SOURCE), as shown in Figure 8-52.

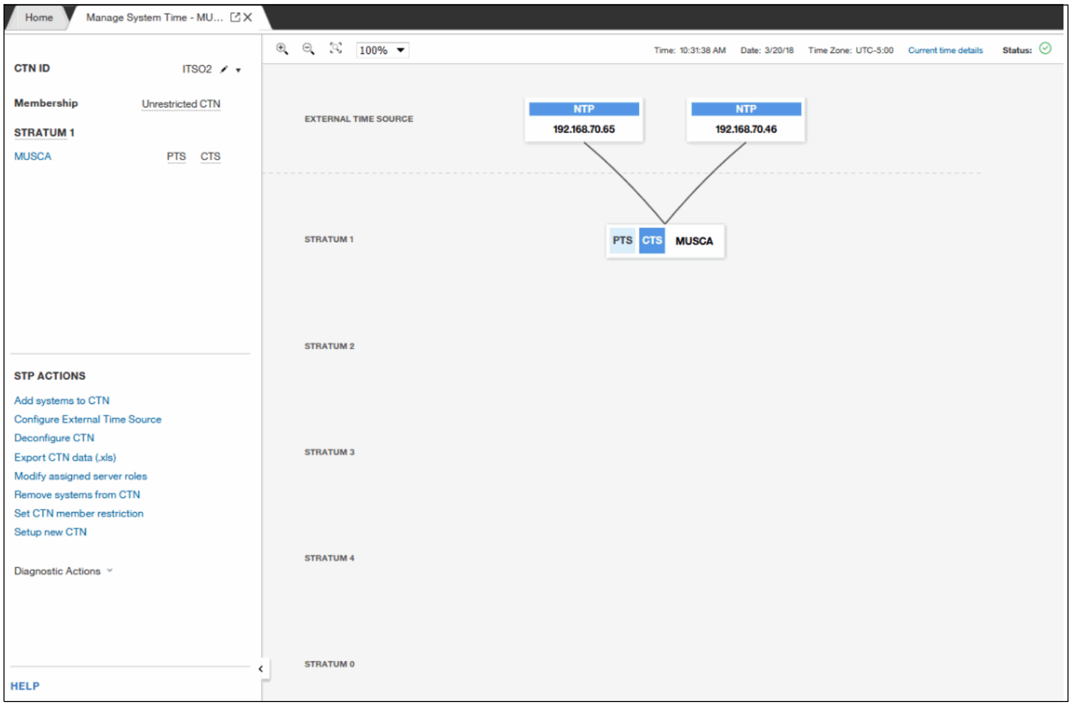

Figure 8-52 Verify ETS configuration

8.7 Adding the z14 ZR1 server to a CTN

This section describes the procedure that is used to add a z14 ZR1 CPC to a CTN. As shown in Figure 8-53, we add a z14 ZR1 (MUSCA) to the CTN with the ID CTLSOMRS.

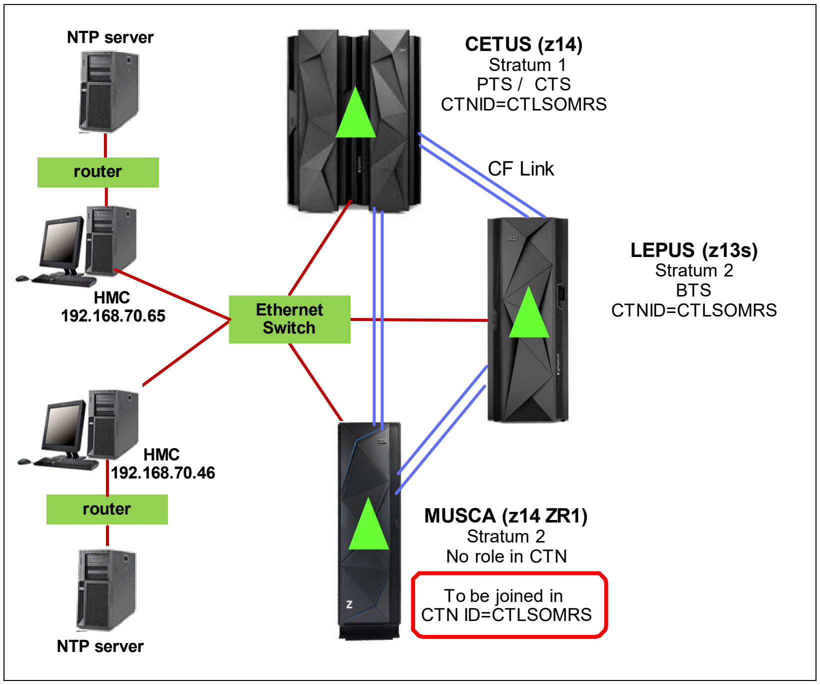

Figure 8-53 Scenario to add a z14 ZR1 CPC to a CTN

At the beginning of the process, the MUSCA CPC is connected by using coupling links to a LEPUS and CETUS, but it is not assigned to any CTN. When added to the CTN, CETUS becomes a Stratum 2 server.

Next, we describe how to assign the Arbiter role to MUSCA.

8.7.1 HMC operations to add the CPC to the CTN

To add a CPC to a STP-only CTC, complete the following steps:

1. Open the Manage System Time task from the HMC, and select **Add systems to CTN** from STP ACTIONS (see Figure 8-54).

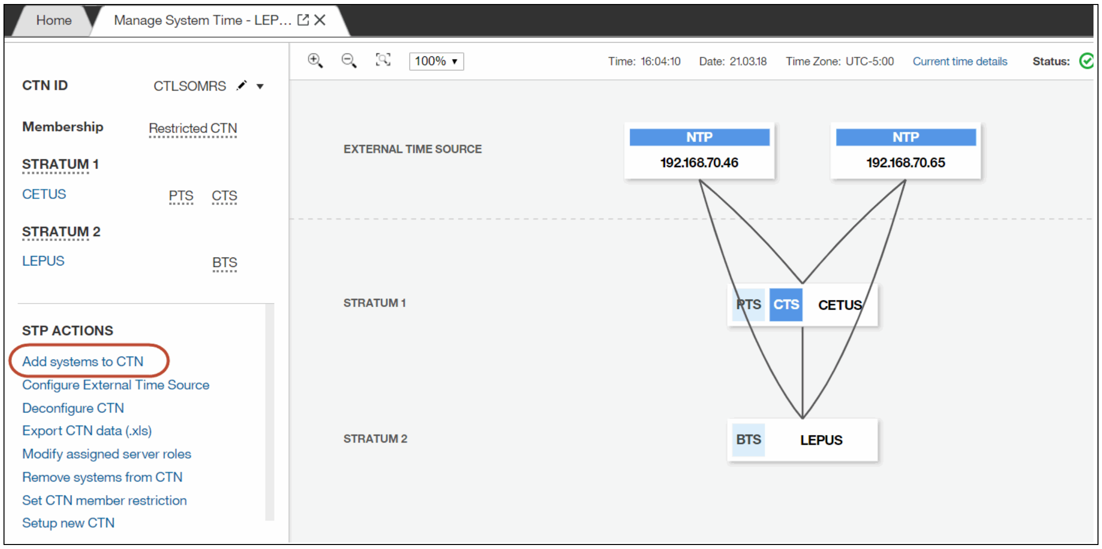

Figure 8-54 Adding systems to CTN

Note: Before a CPC is added to a CTN, ensure that the Allow any server to be a member of the CTN option is set in the Coordinated Timing Network (CTN) member restriction preferences. If the Allow only the servers that are specified below to be members of the CTN option is selected, the “The members of this CTN are restricted” window opens and you cannot add a CPC to CTN (see Figure 8-55).

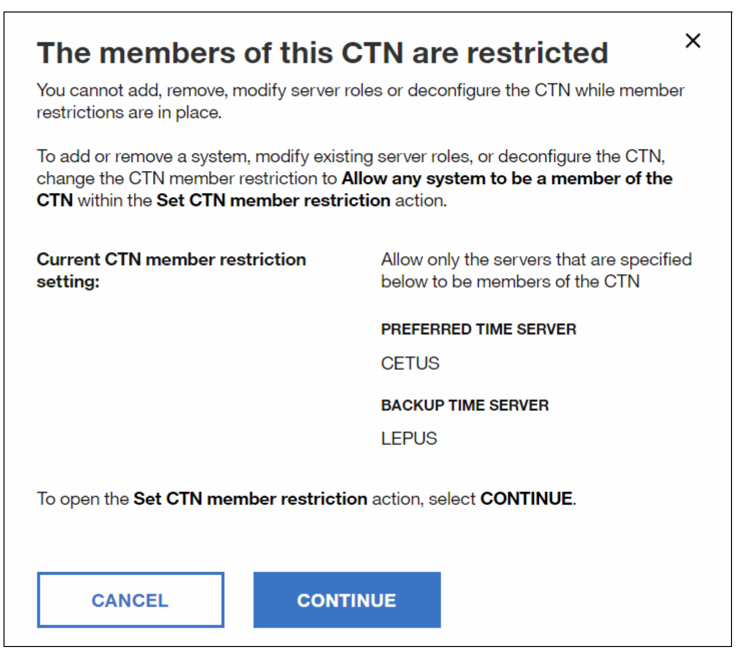

Figure 8-55 The members of this CTN are restricted window

2. The Specify Coordinated Timing Network (CTN) Members window opens. In this window, select the CPC name that does not belong to any CTN (MUSCA). Click **NEXT** to continue (see Figure 8-56).

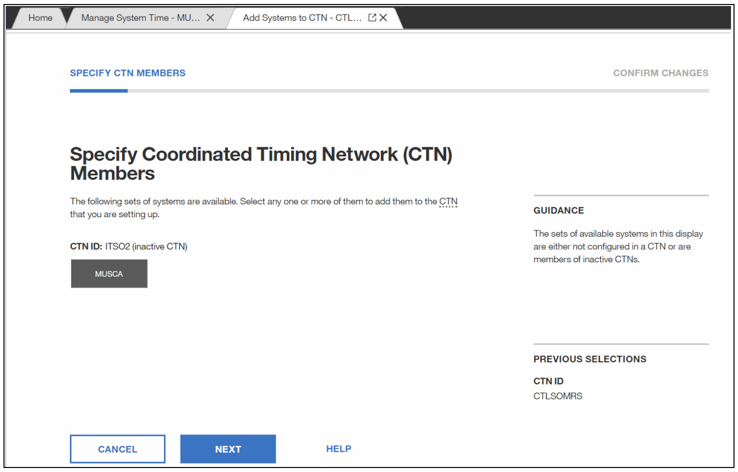

Figure 8-56 Specify Coordinated Timing Network (CTN) Members window

3. The Confirm Changes window opens. In this case, MUSCA is added as a Stratum 2 server. Click **APPLY** to continue (see Figure 8-57).

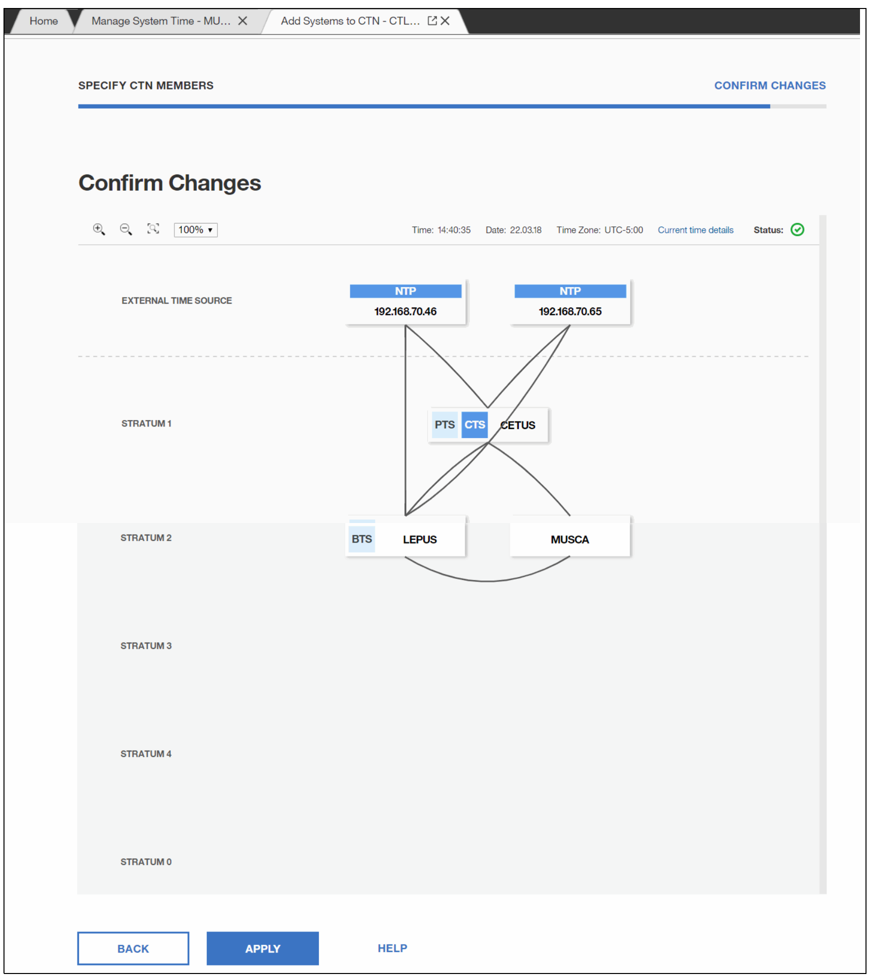

Figure 8-57 Confirm Changes

4. You see the Local CTN ID change confirmation for MUSCA (see Figure 8-58) and after selecting APPLY the Complete message window when finished. Click **Close** (see Figure 8-58).

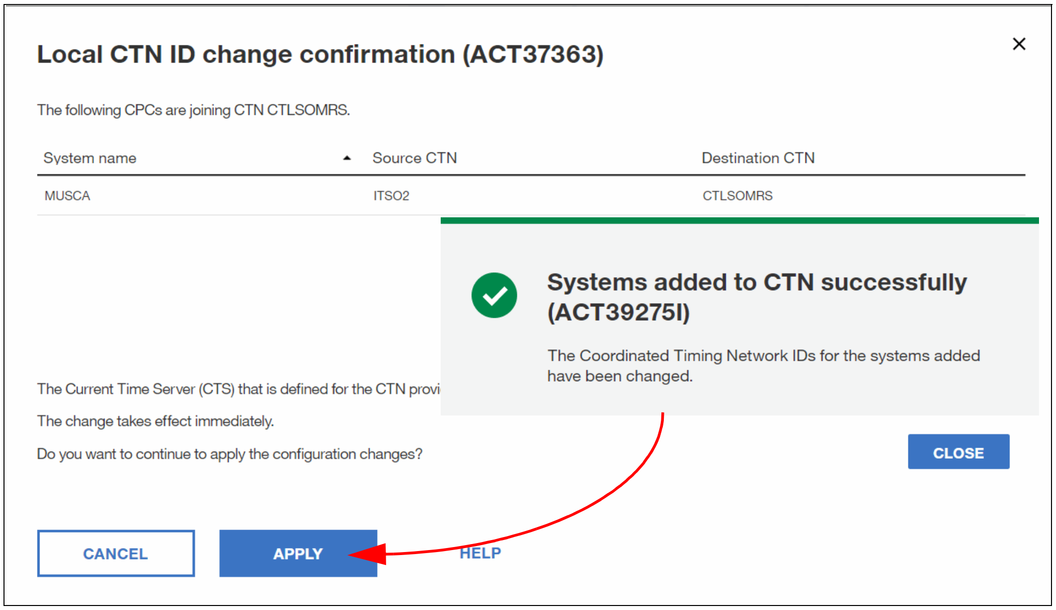

Figure 8-58 Local CTN ID change confirmation (ACT37363) window

8.7.2 Verifying that the system is added to the CTN

From the Manage System Time initial window, check the added CPC's Stratum level and status (see Figure 8-59).

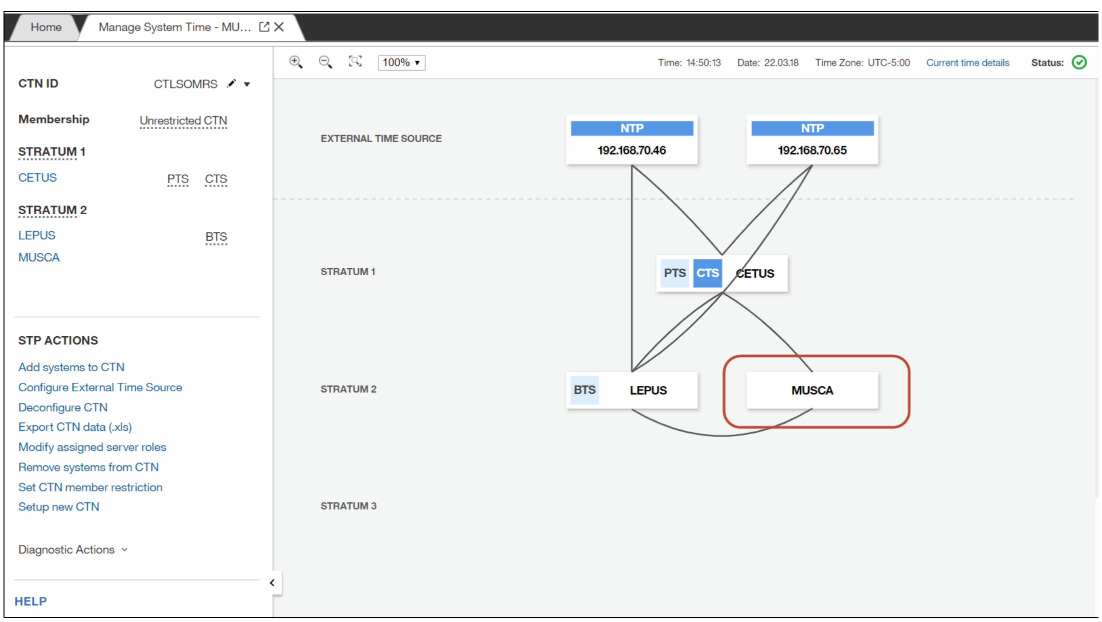

Figure 8-59 After CPC ia added to CTN

8.8 Assign an STP role to the CPC

You can assign an STP role to the CPC in the CTN. PTS must exist, but BTS and Arbiter roles are optional. However, it is a good idea to assign BTS and Arbiter based on your CTN configuration. In this scenario, we assign the Arbiter role to MUSCA, which does not include an STP role (see Figure 8-60).

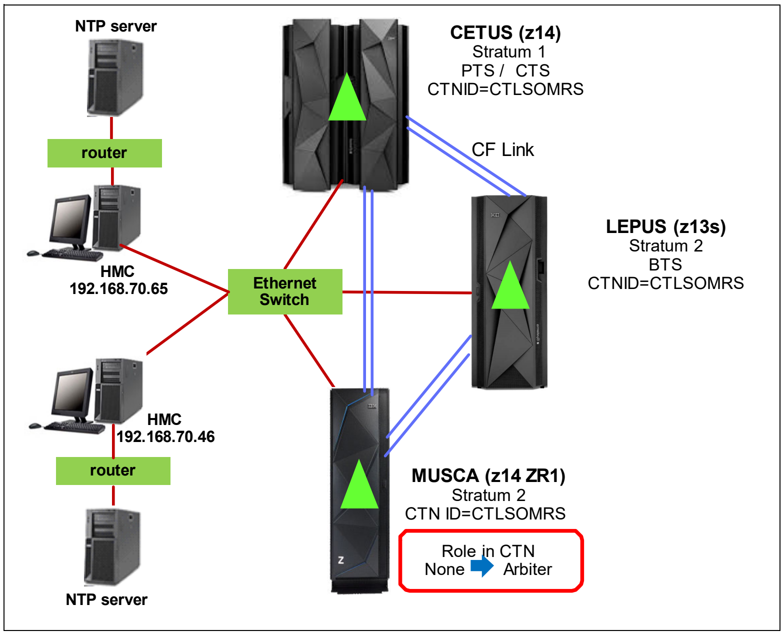

Figure 8-60 Configuration sample that assigns STP role

8.8.1 HMC operations to assign or modify STP Role

To assign or modify STP role, complete the following steps:

1. Open the Manage System Time task from the HMC that can operate CTN changes to modify the STP roles. Select **Modify assigned server roles** from the STP ACTIONS list (see Figure 8-61).

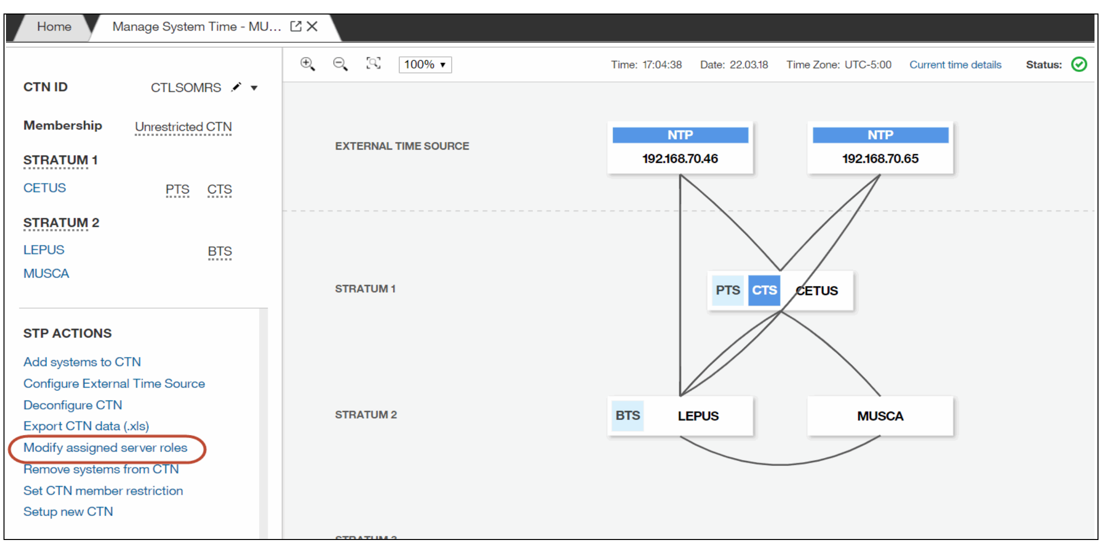

Figure 8-61 Modify server roles

2. The Choose Preferred Time Server window opens. Select the CPC name for the PTS. In this example, we do not change the PTS; therefore, CETUS remains selected (see Figure 8-62).

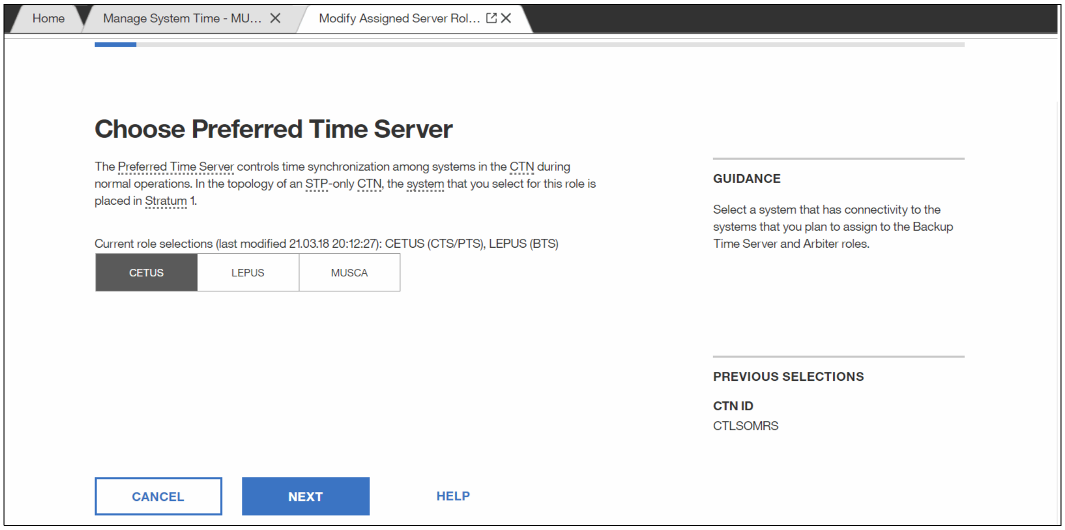

Figure 8-62 Choose Preferred Time Server window

3. The Choose Backup Time Server window opens. In the window, select the CPC name to use for BTS. In this example, we choose LEPUS to remain the BTS. Click **NEXT** to continue (see Figure 8-63).

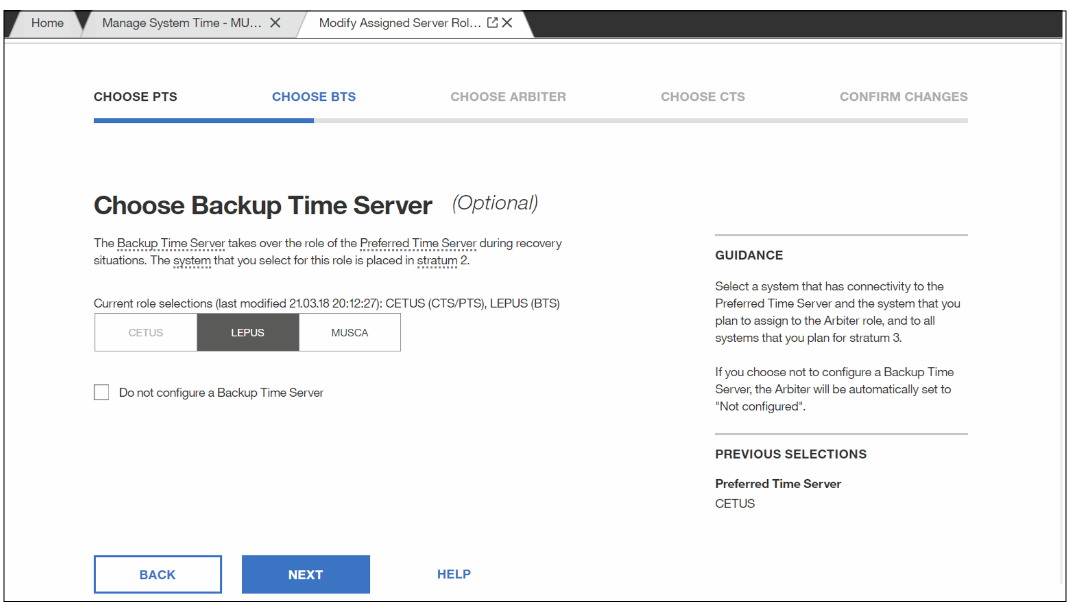

Figure 8-63 Choose Backup Time Server window

4. The Choose Arbiter window opens. For now, a third server is in the CTN; therefore, we can select **MUSCA** to take the Arbiter role. Click **NEXT** to continue (see Figure 8-64).

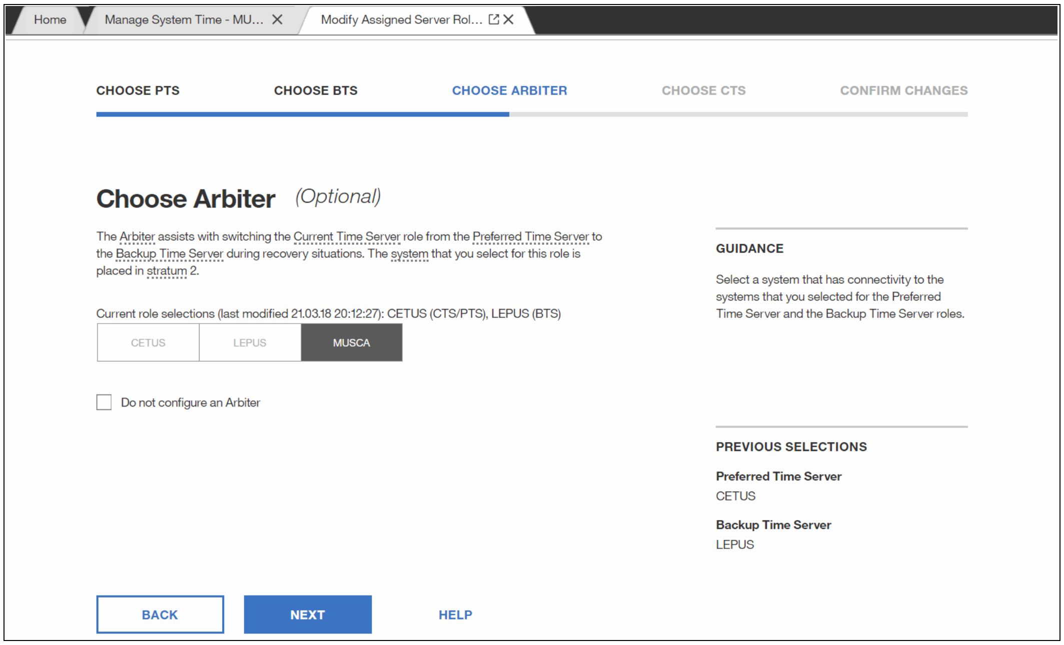

Figure 8-64 Choose Arbiter (Optional) window

5. The Choose Current Time Server window opens. Select the CPC name to act as the CTS. Click **NEXT** to continue (see Figure 8-65).

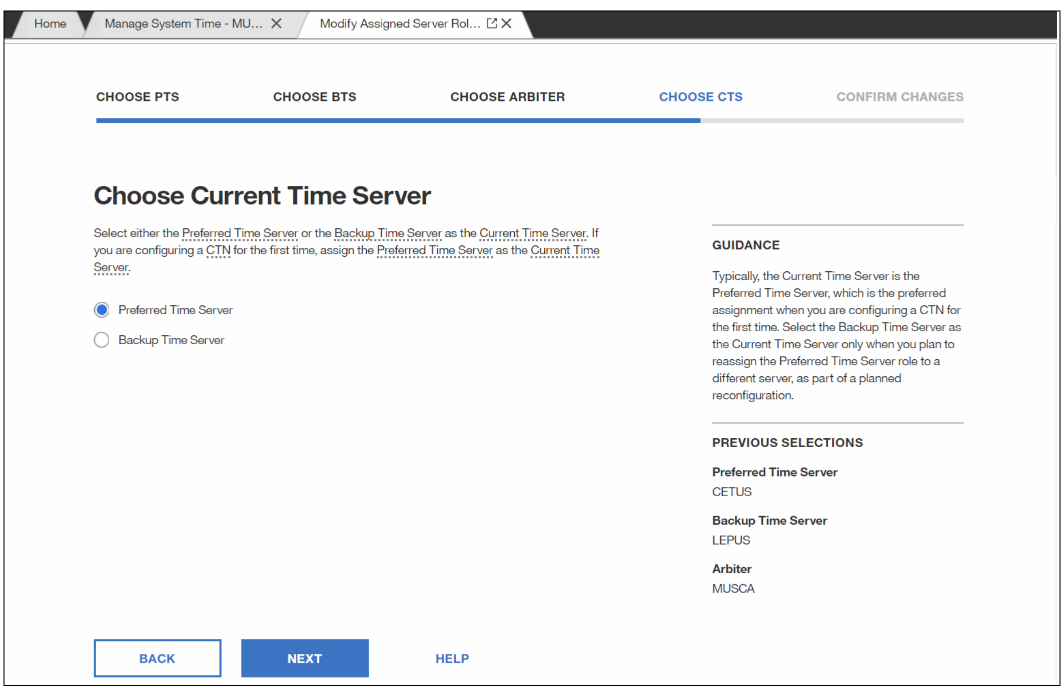

Figure 8-65 Choose Current Time Server window

6. The Confirm Changes window opens. Check the CTN topology. Click **APPLY** to reflect changes (see Figure 8-66).

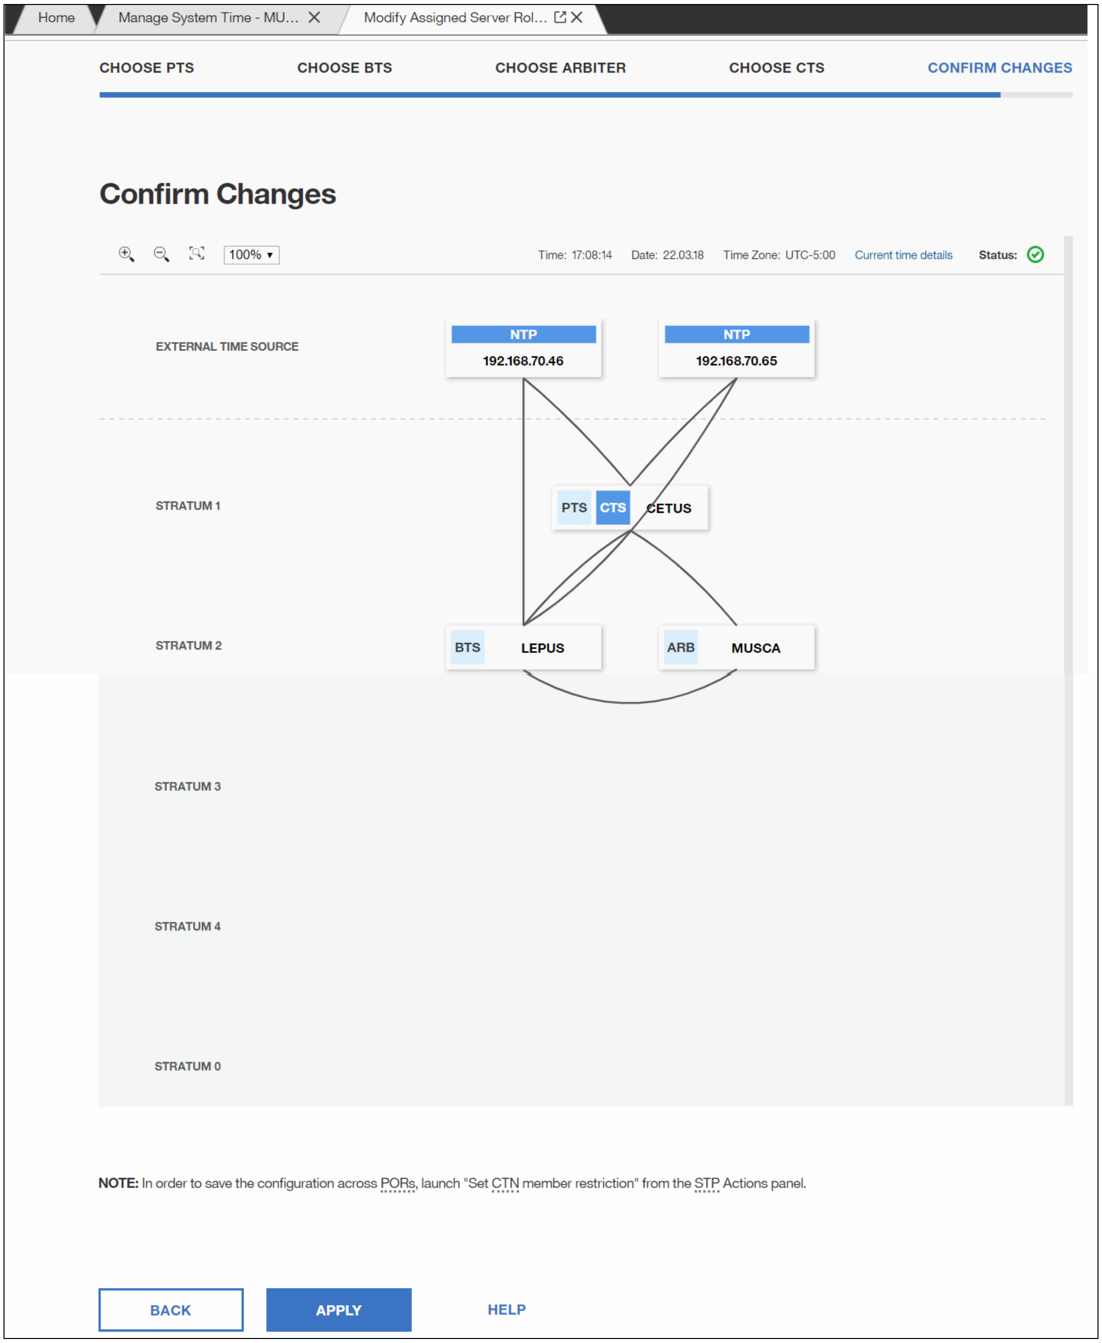

Figure 8-66 Confirm Changes window

7. The Modifying CTN progress is shown until the operation completes successfully, followed by the ACT39295I success window (see Figure 8-67). Click **CLOSE** to return to the initial Manage System Time window.

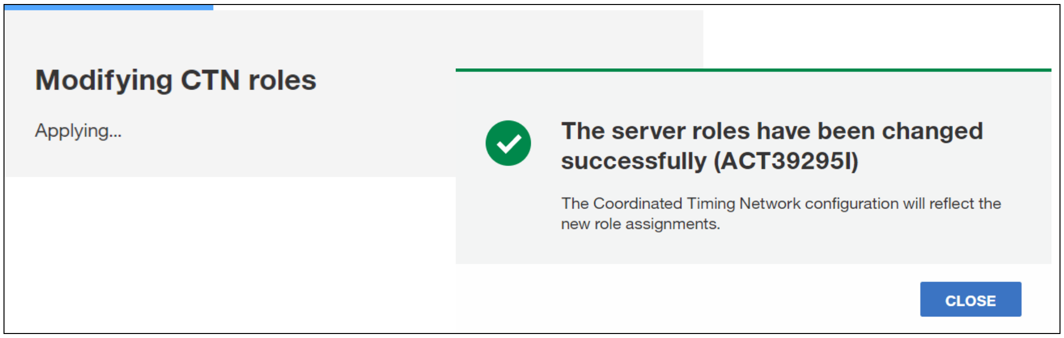

Figure 8-67 Modifying CTN window / ACT392795I Success window

8.8.2 Verifying assigned server roles

In the Manage System Time initial window, check the CPC role assignment. In this example, MUSCA was configured for the Arbiter role (see Figure 8-68).

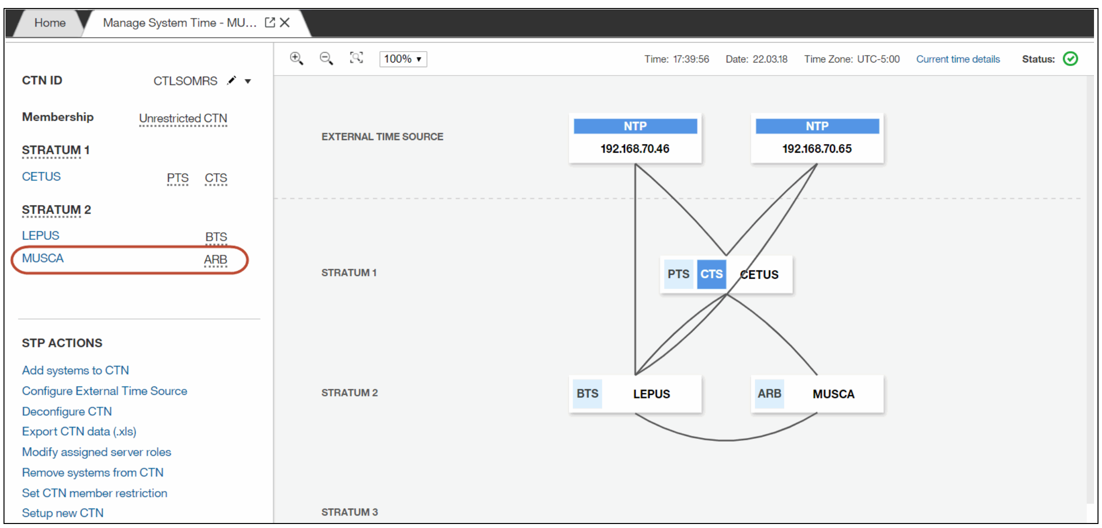

Figure 8-68 Confirm modifying CTN roles

Also, you can confirm STP status by using the z/OS `D ETR` command.

The output of the **D ETR** command before the arbiter role is assigned to MUSCA is shown in Figure 8-69.

```
D ETR
IEA386I 13.06.55 TIMING STATUS 057
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = CTLSOMRS
  THE STRATUM 1 NODE ID = 003906.M03.IBM.02.0000000EE0F7
  NUMBER OF USABLE TIMING LINKS =  7
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

Figure 8-69 D ETR command before LEEPUS is assigned as BTS

The output of the **D ETR** command after the arbiter role is assigned to MUSCA is shown in Figure 8-70.

```
D ETR
IEA386I 13.09.22 TIMING STATUS 064
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = CTLSOMRS
  THE STRATUM 1 NODE ID = 003906.M03.IBM.02.0000000EE0F7
  THIS IS THE ARBITER SERVER
  NUMBER OF USABLE TIMING LINKS =  7
```

Figure 8-70 D ETR command after MUSCA is assigned as Arbiter

9

Defining Coupling Facility links

This chapter describes the coupling connectivity options that are available on z14 ZR1. Coupling links support Parallel Sysplex and STP.

This chapter includes the following topics:

9.1 Coupling connectivity options on z14 ZR1

A Parallel Sysplex (sysplex) is a collection of z/OS systems that cooperate to process workloads by using certain hardware and software products. The hardware and software components that make up a sysplex cooperate to provide higher availability, coordinated systems management, and improved growth potential over a conventional computer system of comparable processing power.

The Server Time Protocol (STP) facility (FC 1021) is required to synchronize the time-of-day (TOD) clocks for systems in a sysplex that run on different CPCs.

For more information about STP configuration, see Chapter 8, “Preparing for Sysplex and configuring Server Time Protocol” on page 159.

A coupling facility (CF) enables parallel processing and improved data sharing for authorized programs that are running in the sysplex. The cross-system extended services (XES) component of z/OS enables applications and subsystems to take advantage of the coupling facility.

In a Parallel Sysplex, the central processor complexes (CPCs) are connected through a coupling facility by using coupling links.

9.1.1 Coupling connectivity for Parallel Sysplex on z14 ZR1

Coupling connectivity for Parallel Sysplex on z14 ZR1 use Coupling Express Long Range (CE LR) and Integrated Coupling Adapter Short Reach (ICA SR). The ICA SR feature is designed to support distances of up to 150 m (492 feet). The CE LR feature supports distances up to 10 km (6.2 miles) unrepeated between systems, and up to 100 km (62 miles) with a qualified Dense Wavelength Division Multiplexer (DWDM). The available options for coupling links on z14 ZR1 are shown in Figure 9-1.

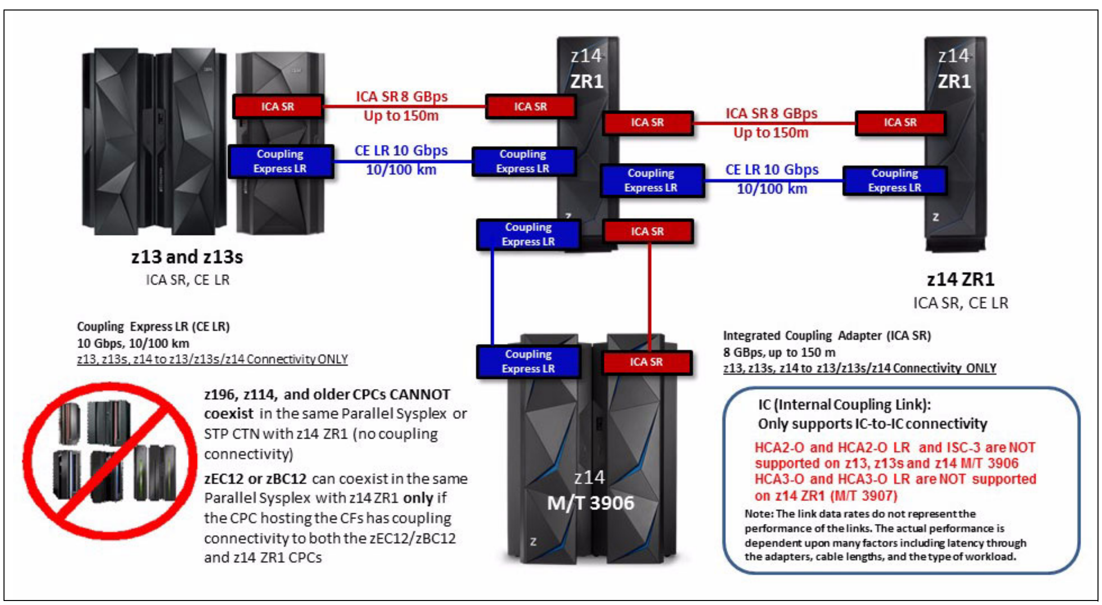

Figure 9-1 z14 ZR1 coupling connectivity

Note: IBM z14 ZR1 does not support HCA3-O fanout for 12x IFB (#0171) and HCA3-O LR fanout for 1x IFB (#0170).

9.1.2 Internal coupling

Internal coupling (IC) links are used for internal communication between LPARs on the same system that is running coupling facilities (CF) and z/OS images. The connection is emulated in Licensed Internal Code (LIC) and provides fast and secure memory-to-memory communications between LPARs within a single system. No physical cabling is required.

9.1.3 Integrated Coupling Adapter Short Reach

First introduced with the IBM z13, the ICA SR is a two-port fanout that is used for short distance coupling connectivity and uses the coupling channel type CS5. The ICA SR also uses PCIe Gen3 technology, with x16 lanes that are bifurcated into x8 lanes for coupling.

The ICA SR supports cable length of up to 150 m (492.1 feet) and supports a link data rate of 8 Gbps[1]. It also supports up to four CHPIDs per port and seven subchannels (devices) per CHPID. The coupling links can be defined as shared between images within a CSS. They also can be spanned across multiple CSSs in an IBM Z CPC.

9.1.4 Coupling Express Long Range

The CE LR is a two-port PCIe native adapter that is used for long-distance coupling connectivity and uses a new coupling channel type: CL5. The CE LR feature uses PCIe Gen3 technology and is hosted in a PCIe I/O drawer. CE LR is not supported in a switched environment (point-to-point links only).

The feature supports communication at unrepeated distances up to 10 km (6.2 miles) by using 9µm single mode fiber optic cables and repeated distances up to 100 km (62 miles) by using IBM Z qualified DWDM vendor equipment. It supports up to four CHPIDs per port and 32 subchannels (devices) per CHPID. CE LR supports a link data rate of up to 10 Gbps[2]. The coupling links can be defined as shared between images within a CSS or spanned across multiple CSSs in an IBM Z CPC.

9.1.5 Preparing to define coupling facility links

A good point to start your preparation for implementing any kind of coupling link is an accurate and current documentation that clearly illustrates all connections that are needed to the new or upgraded CPC.

When installing coupling links, ensure that you ordered enough ports to provide your configuration with physical feature redundancy. Your sysplex should be configured for the highest possible availability.

Sysplex failure independence is a function of a z/OS to CF relationship. For example, all connectors to a structure on a stand-alone CF are failure independent. However, with an ICF, all connections from z/OS images on the same footprint are failure-dependent.

1 The link data rates do not represent the performance of the links. The actual performance depends on many factors, including latency through the adapters, cable lengths, and the type of workload.

2 The link data rates do not represent the performance of the links. The actual performance depends on many factors, including latency through the adapters, cable lengths, and the type of workload.

For more information, see *Coupling Facility Configuration Options*, ZSW01971, which can be found at the Parallel Sysplex on IBM Z web page.

Evaluate whether any channel features in the current configuration are not supported on the z14 ZR1. The configuration is reviewed for any channel types that cannot be carried forward nor connected to the z14 ZR1 server.

Per the earlier Statement for Direction (see 9.1.1, “Coupling connectivity for Parallel Sysplex on z14 ZR1” on page 200), IBM z14 ZR1 server does not support InfiniBand coupling connectivity. Therefore, when deploying a z14 ZR1 CPC, the only coupling links that are supported are the CE LR for long range, and the ICA SR for short reach coupling connectivity.

Another important point to be considered is to ensure that all processors that it will be connected by using coupling links follow the following restrictions: Only IBM Z servers of the current generation (N) or from the previous generation (N-1) can coexist in the same sysplex or CTN if they are assigned to an STP role and use coupling facility structures. The z14 ZR1 can connect to only z14, z13, and z13s if they are equipped with CE LR or ICA SR coupling links.

Note: z14 ZR1 supports coupling connectivity between N and N-1 server generations only. An N-2 server can be attached to an N-1 server for STP timing purposes only.

If coupling links are to be connected cross sites through DWDM, it is necessary to verify whether the DWDM equipment to be used supports the respective couplink link technology and is qualified for Server Time Protocol.

Note: Do *not* use DWDM equipment that is not qualified to transport STP information.

IBM does not provide channel cables as features on the z14 ZR1. Therefore, a complete analysis must be made of the I/O connectors that are used on systems that are being upgraded to z14 ZR1 to ensure that the appropriate fiber cabling is installed.

An equivalent study should be part of your preparation to install a new z14 ZR1 so that all cabling is delivered to the data center before the installation date.

All required cables for the z14 ZR1 must be identified and placed on order. Labeling of all cables it is required for the installation. At a minimum, the labels should identify the PCHID number.

If you received the configuration and PCHID reports from IBM, define your coupling links to fit your planned configuration to your new or upgraded CPC.

Note: Deactivate any coupling link on other connected systems before an upgrade, or you might experience configuration errors.

9.2 Coupling Express Long Range

This section describes implementing couplink links by using the Coupling Express Long reach feature.

9.2.1 CE LR: Implementation

The CE LR coupling link (CHPID type CL5) uses PCIe Gen3 technology and is hosted in a PCIe I/O drawer.

The definition of this CFlink is an example that it is part of the activity Define CF/STP link that is shown in Figure 1-3 on page 5.

A coupling link between a z14 ZR1 CPC and a z14 CPC that uses CL5 CHPIDs is defined in this section as an example (see Figure 9-2).

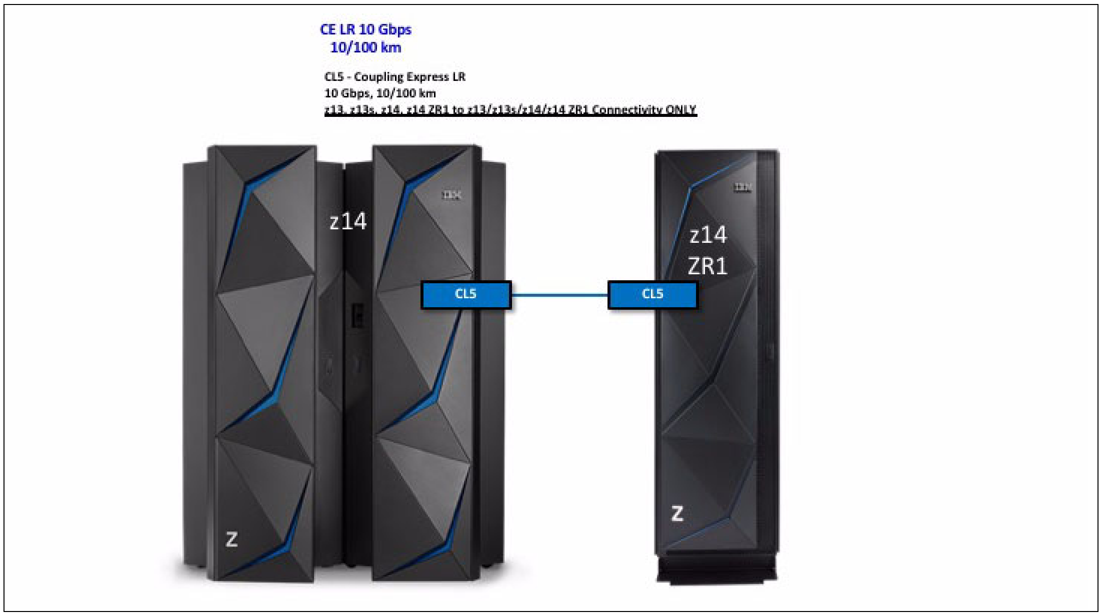

Figure 9-2 CF link connection that is CL5

CE LR is defined in IOCDS similar to PSIFB. Although this feature is a PCIe feature, a PCHID is used instead of an AID to identify the physical card. Example 9-1 shows a sample extract of the IOCP to define the new CHPID Type CL5 on the z14 that connects to a z14 ZR1.

Example 9-1 IOCP definitions for CHPID Type CL5 on z14

```
ID     ..                                                      *
       ..                                                      *
       SYSTEM=(3906,1),LSYSTEM=CETUS,                          *
       TOK=('CETUS',008001117A883907095804670118074F00000000,00*
       000000,'18-03-15','09:58:04','........','........')
RESOURCE PARTITION=((CSS(0),(CETUS0A,A),..                     *
       ),(CETUS0D,D),(CETUS0E,E),(CETUS0F,F),(CETUS01,1),(CETUS*
       ..                                                      *
       (CSS(2),(CETUS2A,A),(CETUS2B,B),(CETUS2C,C),(CETUS2D,D),*
       (CETUS2E,E),..                                          *
CHPID  PATH=(CSS(2),E9),SHARED,PARTITION=((CETUS2E),(=)),      *
       CPATH=(CSS(1),E9),CSYSTEM=MUSCA,PORT=2,PCHID=13C,       *
       TYPE=CL5
```

Example 9-2 on page 204 shows a sample extract of the corresponding IOCP definition for a connecting CHPID Type CL5 on a z14 ZR1.

Example 9-2 IOCP definitions for CHPID Type CL5 on z14

```
ID    ..                                                           *
      ..                                                           *
      SYSTEM=(3907,1),LSYSTEM=MUSCA,                               *
      TOK=('MUSCA',008001117A88390709580467O118074F00000000,00*
      000000,'18-03-15','09:58:04','........','........')
RESOURCE PARTITION=((CSS(0),(MUSCA0A,A),..                         *
      USCA07,7),(MUSCA08,8),(MUSCA09,9)),(CSS(1),(MUSCA1A,A),(*
      MUSCA1B,B),(MUSCA1C,C),(MUSCA1D,D),(MUSCA1E,E),(MUSCA1F,*
      F),(MUSCA11,1),..                                            *
CHPID PATH=(CSS(1),E9),SHARED,PARTITION=((MUSCA11),(=)),           *
      CPATH=(CSS(2),E9),CSYSTEM=CETUS,PORT=2,PCHID=124,            *
      TYPE=CL5
CNTLUNIT CUNUMBR=FFFD,PATH=((CSS(1),E5,E9)),UNIT=CFP
IODEVICE ADDRESS=(FFD4,008),CUNUMBR=(FFFD),UNIT=CFP
IODEVICE ADDRESS=(FFDC,008),CUNUMBR=(FFFD),UNIT=CFP
```

Note: When you connect CF sender and CF receiver channel paths, or CF peer channel paths, HCD proposes coupling facility control unit and device numbers that must be defined for a CF sender channel. (CF receiver channels do not require control units and devices to be defined.)

For more information about how to define CHPID Type CL5 in HCD, see 14.2.4, “Defining CL5 CHPIDs” on page 355 in this book.

9.2.2 CE LR: Managing the configuration

After activating the new configuration with the new CE LR CF links and after all the cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the Support Element (SE) panels:

- Checking status by using z/OS commands:
  - For example, to check the status of CHPID E9, you run a **D M=CHP(E9)** command, as shown in Example 9-3.

Example 9-3 Display status of CHPID E9

```
D M=CHP(E9)
IEE174I 16.04.03 DISPLAY M 738
CHPID E9:  TYPE=34, DESC=COUPLING OVER ROCE, ONLINE
COUPLING FACILITY 003907.IBM.02.000000007A88
                  PARTITION: 1F  CPCID: 00
NAMED CF77        CONTROL UNIT ID: FFFD

PATH        PHYSICAL              LOGICAL  CHANNEL TYPE      CAID PORT
E9 / 0507   ONLINE                ONLINE   CL5 10GbE-RoCE    013C  02

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   48   IN USE:   48   NOT USING:    0   NOT USABLE:    0
 OPERATIONAL DEVICES / SUBCHANNELS:
     FFA4 / 3C40     FFA5 / 3C41     FFA6 / 3C42     FFA7 / 3C43
     FFA8 / 3C44     FFA9 / 3C45     FFAA / 3C46     FFAB / 3C47
     FFBC / 3C48     FFBD / 3C49     FFBE / 3C4A     FFBF / 3C4B
     FFC0 / 3C4C     FFC1 / 3C4D     FFC2 / 3C4E     FFC3 / 3C4F
```

```
FFC4 / 3C50     FFC5 / 3C51     FFC6 / 3C52     FFC7 / 3C53
FFC8 / 3C54     FFC9 / 3C55     FFCA / 3C56     FFCB / 3C57
FFCC / 3C58     FFCD / 3C59     FFCE / 3C5A     FFCF / 3C5B
FFD0 / 3C5C     FFD1 / 3C5D     FFD2 / 3C5E     FFD3 / 3C5F
FFD4 / 3C60     FFD5 / 3C61     FFD6 / 3C62     FFD7 / 3C63
FFD8 / 3C64     FFD9 / 3C65     FFDA / 3C66     FFDB / 3C67
FFDC / 3C68     FFDD / 3C69     FFDE / 3C6A     FFDF / 3C6B
FFE0 / 3C6C     FFE1 / 3C6D     FFE2 / 3C6E     FFE3 / 3C6F
```

- Checking status by using SE panels:
  a. From the HMC, select the CEC (under Systems Management) where the CHPID/PCHID you want to verify is stored, and click **Single Object Operations** (under Recovery task options).
  b. On the SE, select the same CEC and click **Channels**, as shown in Figure 9-3.

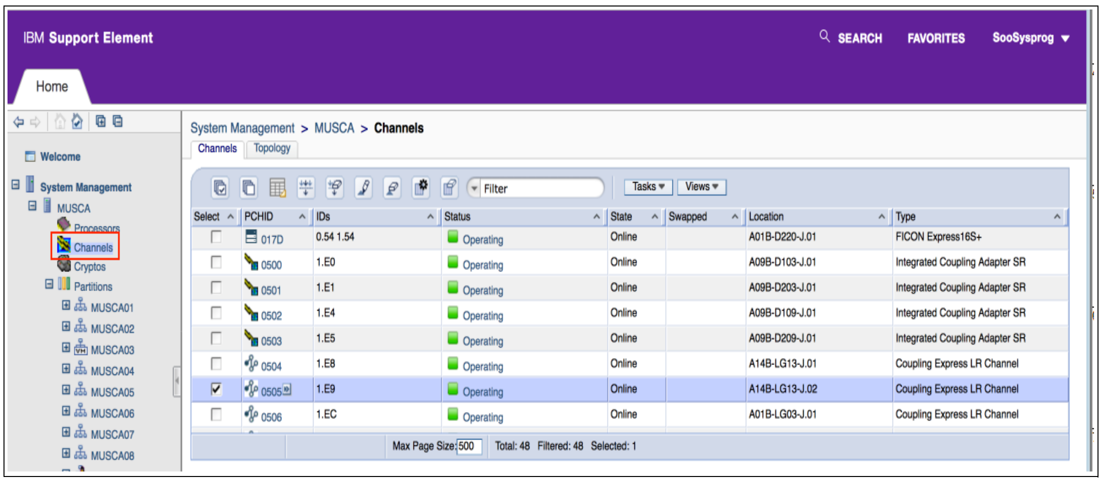

Figure 9-3 SE Systems Management and channels

  c. Look for the PCHID of which you are interested in checking the status. The result resembles the result that is shown in Figure 9-4.

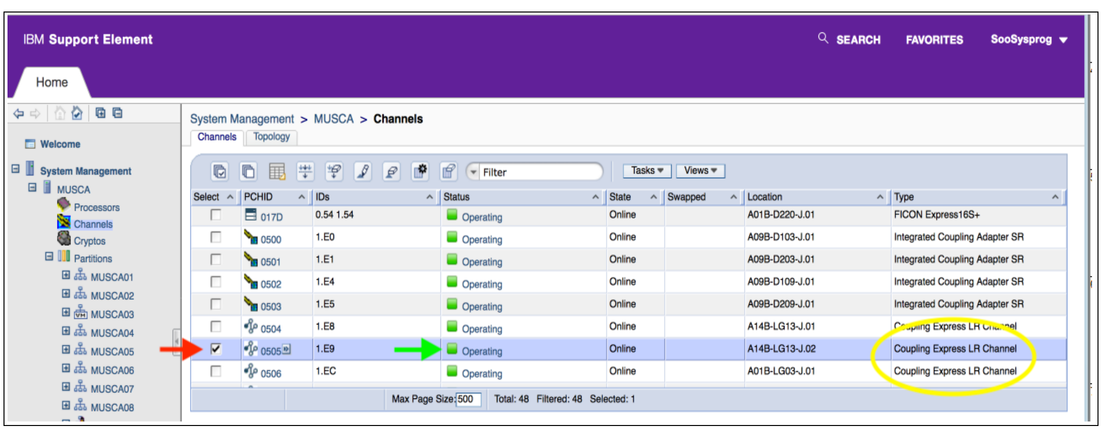

Figure 9-4 Verifying channel CL5 by using CEC view

  d. For more information about the PCHID, click the PCHID to show the details, as shown in Figure 9-5 on page 206.

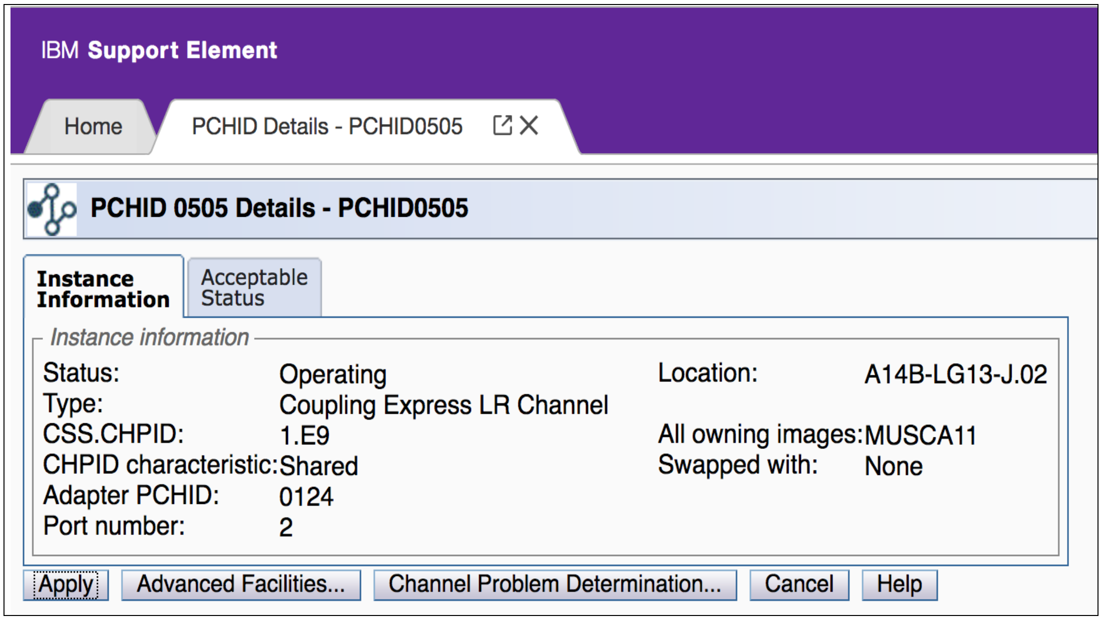

Figure 9-5 CL5 PCHID details

You can also check the status from the LPAR view by selecting the LPAR that you want to check the CHPID status. Then, select the channels option under that LPAR. Now you can search for the CHPID and check the status, as shown in Figure 9-6.

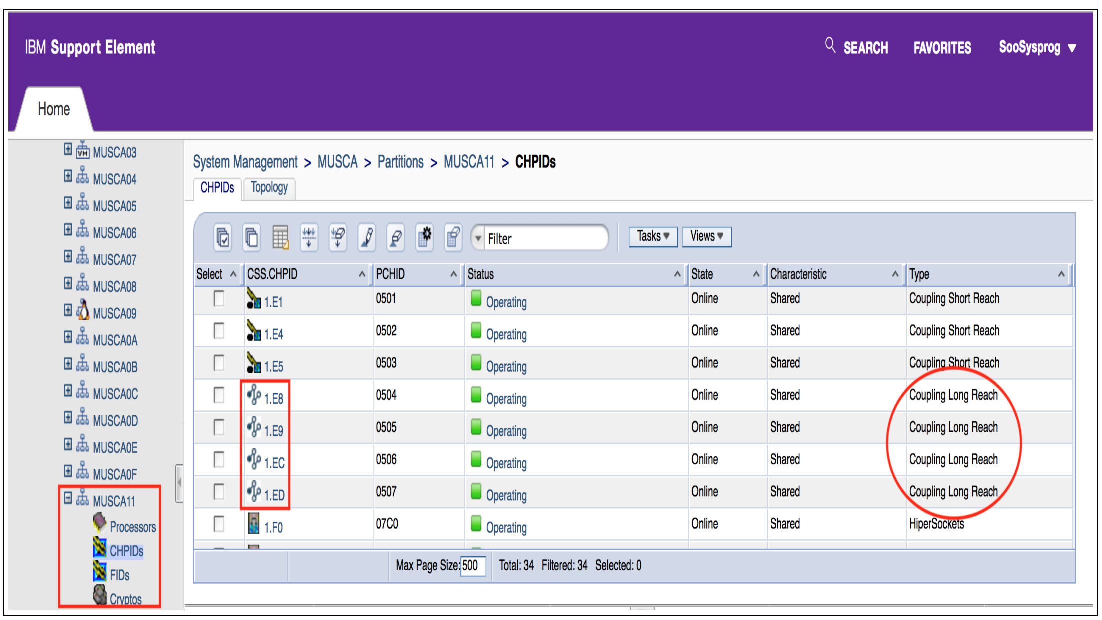

Figure 9-6 SE Verify channel LPAR view

9.3 Integrated Coupling Adapter Short Range

This section describes implementing couplink links by using the Integrated Coupling Adapter Short Reach (ICA SR) feature.

9.3.1 ICA SR: Implementation

The CHPID type (CS5) was introduced with the z13 and uses a PCIe-Gen3 fanout feature that is named Integrated Coupling Adapter FC 0172.

The definition of this CF link is an example that it is part of the activity Define CF/STP link, as shown in Figure 1-3 on page 5.

A coupling link connection between a z14 ZR1 and a z14 that uses CS5 CHPIDs is shown in Figure 9-7.

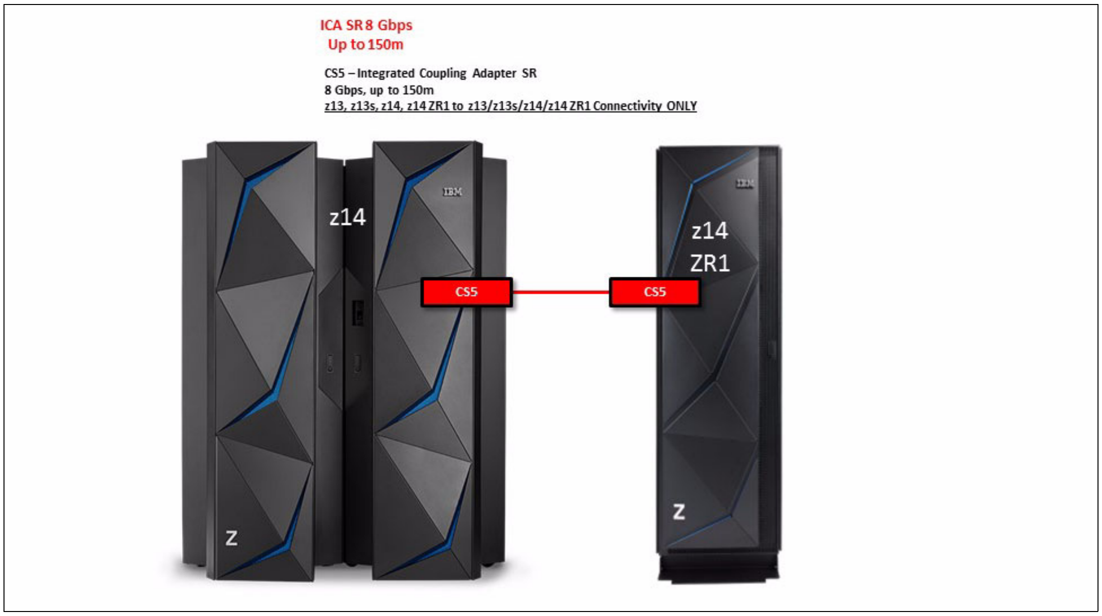

Figure 9-7 CF link CS5 connection from z14 ZR1 to z14

The ICA SR is defined in IOCDS similar to PSIFB, by using an AID to identify the physical card. Example 9-4 shows a sample of the IOCP that is defining the CS5 CHPID.

Example 9-4 IOCP definitions for CHPID Type CS5 on z14

```
ID    ..                                                        *
      ..                                                        *
      SYSTEM=(3906,1),LSYSTEM=CETUS,                            *
      TOK=('CETUS',008001117A88390709580467011807 4F00000000,00*
      000000,'18-03-15','09:58:04','........','........')
RESOURCE PARTITION=((CSS(0),(CETUS0A,A),..                      *
      (CSS(2),(CETUS2A,A),(CETUS2B,B),(CETUS2C,C),(CETUS2D,D),*
      (CETUS2E,E),..                                            *
CHPID PATH=(CSS(2),E0),SHARED,PARTITION=((CETUS2E),(=)),       *
      CPATH=(CSS(1),E5),CSYSTEM=MUSCA,AID=20,PORT=1,TYPE=CS5
```

Example 9-5 shows a sample extract of the corresponding IOCP definition for a connecting CHPID Type CS5 on a z14 ZR1.

Example 9-5 IOCP definitions for CHPID Type CS5 on z14

```
ID    ..                                                    *
      ..                                                    *
      SYSTEM=(3907,1),LSYSTEM=MUSCA,                        *
      TOK=('MUSCA',008001117A88390709580467011807 4F00000000,00*
      000000,'18-03-15','09:58:04','........','........')
RESOURCE PARTITION=((CSS(0),(MUSCA0A,A),..                  *
      USCA07,7),(MUSCA08,8),(MUSCA09,9)),(CSS(1),(MUSCA1A,A),(*
      MUSCA1B,B),(MUSCA1C,C),(MUSCA1D,D),(MUSCA1E,E),(MUSCA1F,*
      F),(MUSCA11,1),..                                     *
CHPID PATH=(CSS(1),E5),SHARED,PARTITION=((MUSCA11),(=)),    *
      CPATH=(CSS(2),E0),CSYSTEM=CETUS,AID=16,PORT=2,TYPE=CS5
CNTLUNIT CUNUMBR=FFFD,PATH=((CSS(1),E5,E9)),UNIT=CFP
IODEVICE ADDRESS=(FFD4,008),CUNUMBR=(FFFD),UNIT=CFP
IODEVICE ADDRESS=(FFDC,008),CUNUMBR=(FFFD),UNIT=CFP
```

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 352.

9.3.2 ICA SR: Managing 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 z/OS or the SE panels:

- Checking status by using the z/OS command:
  - For example, to check the status of CHPID E1, run a **D M=CHP(E1)** command, as shown in Example 9-6.

Example 9-6 Display status of CHPID E1

```
D M=CHP(E1)
IEE174I 16.20.22 DISPLAY M 758
CHPID E1:  TYPE=33, DESC=COUPLING OVER PCIE, ONLINE
COUPLING FACILITY 003907.IBM.02.000000007A88
                  PARTITION: 1F  CPCID: 00
NAMED CF77        CONTROL UNIT ID: FFFD

PATH        PHYSICAL              LOGICAL  CHANNEL TYPE      CAID PORT
E1 / 0503   ONLINE                ONLINE   CS5 8X-PCIE3      0020  02

COUPLING FACILITY SUBCHANNEL STATUS
 TOTAL:   48   IN USE:   48   NOT USING:    0   NOT USABLE:    0
 OPERATIONAL DEVICES / SUBCHANNELS:
     FFA4 / 3C40     FFA5 / 3C41     FFA6 / 3C42     FFA7 / 3C43
     FFA8 / 3C44     FFA9 / 3C45     FFAA / 3C46     FFAB / 3C47
     FFBC / 3C48     FFBD / 3C49     FFBE / 3C4A     FFBF / 3C4B
     FFC0 / 3C4C     FFC1 / 3C4D     FFC2 / 3C4E     FFC3 / 3C4F
     FFC4 / 3C50     FFC5 / 3C51     FFC6 / 3C52     FFC7 / 3C53
     FFC8 / 3C54     FFC9 / 3C55     FFCA / 3C56     FFCB / 3C57
     FFCC / 3C58     FFCD / 3C59     FFCE / 3C5A     FFCF / 3C5B
     FFD0 / 3C5C     FFD1 / 3C5D     FFD2 / 3C5E     FFD3 / 3C5F
     FFD4 / 3C60     FFD5 / 3C61     FFD6 / 3C62     FFD7 / 3C63
```

```
FFD8 / 3C64     FFD9 / 3C65     FFDA / 3C66     FFDB / 3C67
FFDC / 3C68     FFDD / 3C69     FFDE / 3C6A     FFDF / 3C6B
FFE0 / 3C6C     FFE1 / 3C6D     FFE2 / 3C6E     FFE3 / 3C6F
```

- Checking status by using SE panels:

  a. From the HMC, select the CEC (under Systems Management) where is the CHPID/PCHID you want to verify, and click **Single Object Operations** (under Recovery task options).

  b. On the SE, select the same CEC and click **Channels**, as shown in Figure 9-8.

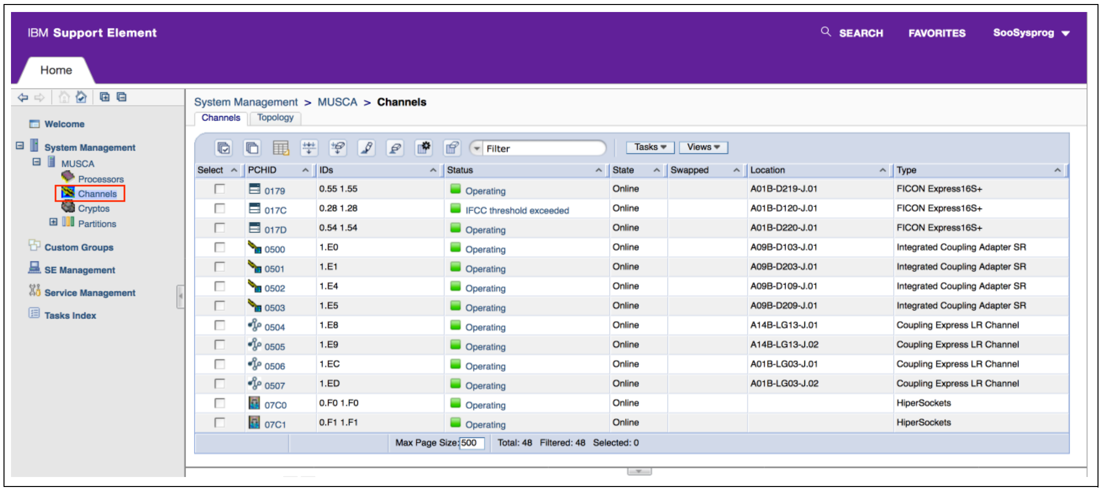

Figure 9-8 SE Systems Management and channels

  c. Look for the PCHID for which you are interested in checking the status. The result resembles the result that is shown in Figure 9-9.

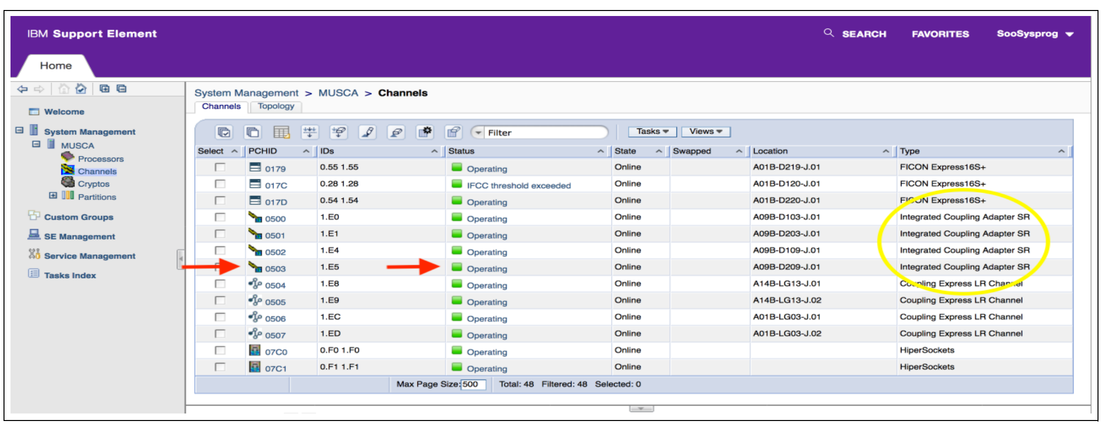

Figure 9-9 Verifying channel CS5 using CEC view

d. For more information about the PCHID, click the PCHID to show a result (see Figure 9-10).

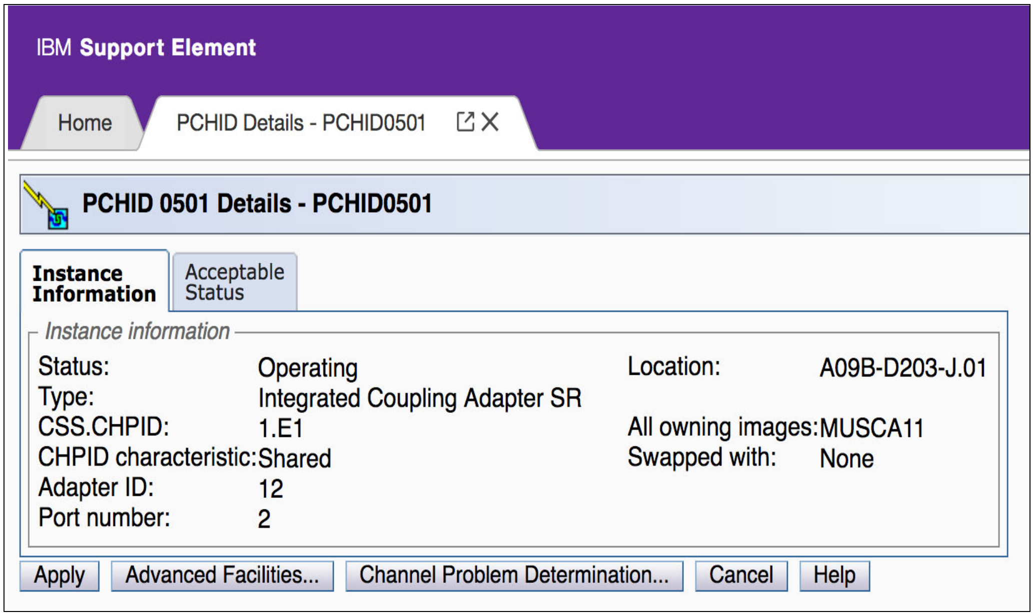

Figure 9-10 CS5 PCHID details

Another option is to check the status from the LPAR view by selecting the LPAR for which you want to check the CHPID status, and select the **Channels** option under that LPAR. Now, you can search for the CHPID and check the status, as shown in Figure 9-11.

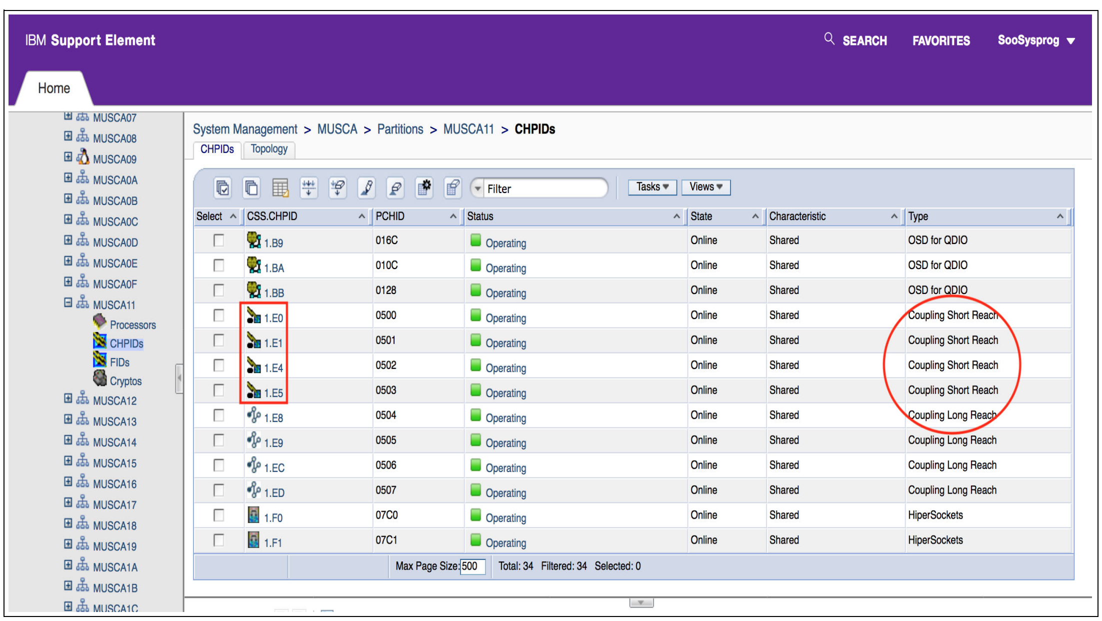

Figure 9-11 SE Verify channel LPAR view

9.4 Defining an STP timing-only link by using ICA SR

This section describes how to configure timing-links (for STP messages) over ICA SR connectivity.

9.4.1 STP timing only: Implementation

In this section, a CHPID type (CS5) connection is used to show an example of 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 a z14 ZR1 and a z14 that uses CS5 CHPIDs is used as an example of how to define an STP timing-only links (see Figure 9-12).

Figure 9-12 STP timing only link connection from z14 ZR1 to z14 (over CS5)

Timing-only link ICA SR is defined in IOCDS similar to ICA SR. The only difference is the control unit type, which is STP in this case. Example 9-7 shows a sample IOCP that defines the CS5 CHPID for timing-only links.

Example 9-7 IOCP defining STP timing only link on a z14 using CS5

```
ID     ..                                                     *
       ..                                                     *
       SYSTEM=(3906,1),LSYSTEM=CETUS,                         *
       TOK=('CETUS',00800112E0F7390609470424011806 1F00000000,00*
       000000,'18-03-02','09:47:04','SYS9','IODF78')
RESOURCE PARTITION=((CSS(0),(CETUS0A,A),..                    *
       (CSS(2),(CETUS2A,A),(CETUS2B,B),(CETUS2C,C),(CETUS2D,D),*
       (CETUS2E,E),(CETUS2F,F),(CETUS21,1),..                 *
CHPID PATH=(CSS(2),E0),SHARED,PARTITION=((CETUS21),(=)),      *
       CPATH=(CSS(1),E5),CSYSTEM=MUSCA,AID=20,PORT=1,TYPE=CS5
CHPID PATH=(CSS(2),E1),SHARED,PARTITION=((CETUS21),(=)),      *
       CPATH=(CSS(1),E0),CSYSTEM=MUSCA,AID=20,PORT=2,TYPE=CS5
CNTLUNIT CUNUMBR=FFFD,PATH=((CSS(2),E9,ED,E1,E4,E5,E0)),      *
       UNIT=STP
```

The sequence of steps to define the STP timing only link connection between a z14 ZR1 and a z14 that uses CS5 CHPIDs by using HCD is the same as for defining the CF links to a connection by using CS5 CHPIDs. Complete the procedure as described in 9.3.2, “ICA SR: Managing the configuration” on page 208. Consider the following points:

- On the CF links connection step, after including the data that is related to the CPC CETUS side of the connection, enter `YES` on the Timing-only link option of the Connect to CF Channel Path panel (see Figure 9-13) and press Enter.

```
  Goto  Filter  Backup  Query  Help
- *---------------- Connect to CF Channel Path -----------------*
  |                                                             |
C |                                                             |
  | Specify the following values.                               |
S |                                                             |
  | Source processor ID  . . . . . : MUSCA                      |
S | Source channel subsystem ID  . : 1                          |
S | Source channel path ID . . . . : E0                         |
S | Source channel path type . . . : CS5                        |
  |                                                             |
  | Destination processor ID . . . . . . CETUS     +            |
/ | Destination channel subsystem ID . . 2   +                  |
p | Destination channel path ID  . . . . E1  +                  |
_ |                                                             |
_ | Timing-only link . . . . . . . . . . YES                    |
_ |                                                             |
_ |                                                             |
_ |                                                             |
_ |                                                             |
_ *-------------------------------------------------------------*
```

Figure 9-13 STP timing only link from z14 ZR1 to z13 by using CS5

- After accepting or overriding the Control Unit and Device numbers for both processors, HCD returns to the CF Channel Path Connectivity List panel (see Figure 9-14). You can see that only the STP timing links are now connected.

```
                     CF Channel Path Connectivity List            Row 1 of
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 1
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------  -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode  Type  Dev
_ E0  12/1  N  CS5  SHR  N     CETUS.2    E1  20/2  N  CS5  SHR   STP
```

Figure 9-14 Only STP Timing links connected

9.4.2 STP timing-only links: Managing the configuration

After activating the new configuration with the ICA SR links that are defined as STP timing-only links and after all cables are connected, verify whether the CHPIDs are online and operating by using z/OS or the SE panels.

The same process that you used for ICA SR links is 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 9.3.2, “ICA SR: Managing the configuration” on page 208 by using the respective CHPID/PCHID of the links you are interested in checking the status.

9.5 CF LPAR setup and CFCC Level 22

In this section, we remind you of some configuration aspects that are related to a Coupling Facility LPAR, such as THIN INTERRUPT, MODE, memory considerations because of the new CF Level, and the changes in CFCC Level 22.

As shown in Figure 1-3 on page 5, the following aspects must be considered for the activity Setup CF LPAR:

- THIN INTERRUPT

  The use of this enhancement can improve the responsiveness of shared ICFs or CPs that are running CFCC in a test or development environment.

  However, production environments often use dedicated ICFs or dedicated CPs.

- MODE

  The default setting for a 3907 Coupling Facility Partition is Volatile. Ensure that plans are in place to set this setting to Non-volatile, if required.

- Memory considerations

  Memory planning must consider the CFCC memory and structure size increases that are associated with a new level of the Coupling Facility Control Code (CFCC).

  LPARS that are running the CFCC code include increased storage requirements:

  – CF Level 21 (or earlier) to CF Level 22
  – CFSizer Tool

  Also, as in previous CF Levels, ensure that the CF LPAR includes at least 512 MB of storage for CFCC code.

9.5.1 Coupling Facility Level 22

The new CFCC Level 22 introduces changes and improvements in the following areas:

- Coupling Facility processor scalability

  Coupling Facility work management and dispatcher changes to allow improved efficiency as processors are added to scale up the capacity of a CF image:

  – Non-prioritized (FIFO-based) work queues

  – Master/subordinate system-managed duplexing protocol Processor scalability for CF images with five or more dedicated processors

  – Use of improved instruction set that was introduced with z14

  These changes automatically apply to all CF images on z14 ZR1

- XCF/XES List Notification Enhancements:

  – CF list structures support three notification mechanisms to inform users about the status of shared objects in the CF:

    • List: This information is used by many users, including XCF Signaling
    • Key-range: Used predominantly by IBM WebSphere® MQ shared queues
    • Sublist notification: Used predominantly by IMS shared queues

– CF users requested the following enhancements to these notification mechanisms:
  - Immediate/delayed round-robin notification for list and key-range notifications (requested by WebSphere MQ)
  - Aggressive list and key-range notifications (requested by IBM WebSphere MQ)
  - List full/not-full notifications (requested by XCF Signaling)

10

Specialized features

This chapter describes the configuration of all specialized features that are available for IBM z14 ZR1. It includes the following topics:

10.1 Crypto Express6S

In this section, we describe the configuration of the Crypto Express6S feature on an IBM z14 ZR1. We also cover cryptographic domains, configuration rules, and what to consider when you are planning for nondisruptive installation of cryptographic features.

10.1.1 Crypto Express6S overview

The following generations of cryptographic coprocessors[1] are supported in z14 ZR1:

- Crypto Express5S, Feature Code #0890, carry forward only (MES from z13s)
- Crypto Express6S, Feature Code #0893

This section describes the Crypto Express6S feature, with the new z14 ZR1. For more information about Crypto Express5S feature configuration (not changed from previous server generation [z13]), see *IBM z13 Configuration Setup*, SG24-8260.

Each of cryptographic coprocessor of z14 ZR1 includes 40 physical sets of registers, which correspond to the maximum number of LPARs running on a z14 ZR1. Each of these 40 sets belongs to the following domains:

- A cryptographic domain index, in the range of 0 - 39, is allocated to a logical partition by the definition of the partition in its image profile. The same domain must also be allocated to the ICSF instance running in the logical partition that uses the Options data set.
- Each ICSF instance accesses only the Master Keys that are corresponding to the domain number that is specified in the logical partition image profile at the Support Element 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 are allocated to this logical partition.

The installation of the CP Assist for Cryptographic Functions (CPACF) DES/TDES enablement (FC 3863) is one of the prerequisites for the use of the special hardware cryptographic feature in z14 ZR1. This feature enables the following functions:

- For data privacy and confidentially: Data Encryption Standard (DES) includes single-length key DES, double-length key DES, and triple-length key DES (also known as TDES). It also includes Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys.
- For data integrity:
  - Secure Hash Algorithm-1 (SHA-1) for 160-bit support
  - SHA-2 for 224-bit, 256-bit, 384-bit, and 512-bit support
  - SHA-3 for 224-bit, 256-bit, 384-bit, and 512-bit support
  - SHAKE for 128-bit and 256-bit support

  SHA-1, SHA-2, and SHA-3 are included enabled on all IBM z14 ZR1 and do not require the no-charge enablement FC 3863.
- For Key Generation: Pseudo-Random Number Generation (PRNG), Deterministic Random Number Generation (DRNG), and True Random number generation (TRNG).
- For message authentication code (MAC): Single-key MAC and double-key MAC.

The total number of Crypto Express6S or Crypto Express5S features cannot exceed 16 per z14 ZR1. The initial order for Crypto Express6S is two features (two PCIe adapters for Crypto Express6S). After the initial order, the minimum order is one feature.

[1] Regional Crypto Enablement is also supported, such as on z13, but this topic is not covered here.

Each Crypto Express6S feature contains one PCIe adapter. The adapter can be in the following configurations:

- Common Cryptographic Architecture (CCA) Coprocessor (CEX6C)
- Public Key Cryptography Standards (PKCS) #11 (EP11) Coprocessor (CEX6P)
- Accelerator (CEX6A)

During the feature installation, the PCIe adapter is configured by default as the CCA coprocessor.

The configuration of the Crypto Express6S adapter as EP11 coprocessor requires a Trusted Key Entry (TKE) workstation Hardware 9.0 (FC 0085 for the rack-mounted workstation, FC 0086 for the tower workstation) with TKE 9.0 Licensed Internal Code (FC 0879).

The Crypto Express6S feature does not use CHPIDs from the channel subsystem pool. However, the Crypto Express6S feature requires one slot in a PCIe I/O drawer, and one physical channel ID (PCHID) for each PCIe cryptographic adapter.

The cryptographic feature codes for z14 ZR1 are listed in Table 10-1.

Table 10-1 Cryptographic feature codes

Feature code	Description
3863	CP Assist for Cryptographic Functions (CPACF) enablement: This feature is a prerequisite to use CPACF (except for SHA-1, SHA-2, and SHA-3) and cryptographic coprocessor hardware.
0893	Crypto Express6S card: A maximum of 16 features can be ordered (minimum of two adapters). This feature is optional; each feature of which contains one PCI Express cryptographic adapter (adjunct processor). This feature is supported in z14 and z14 ZR1 only.
0890	Crypto Express5S card: This feature cannot be ordered for a new build z14 ZR1, but only on a carry forward MES from z13s. The maximum supported number of Crypto Express5S and Crypto Express6S is 16 in total. This feature is optional; each feature of which contains one PCI Express cryptographic adapter (adjunct processor). This feature is supported in z14, z14 ZR1, z13, and z13s servers only.
0086	Trusted Key Entry (TKE) tower workstation: A TKE provides basic key management (key identification, exchange, separation, update, and backup) and security administration. It is optional for running a Crypto Express6S card in CCA mode and required for running it in EP11 mode. The TKE workstation includes one Ethernet port, and supports connectivity to an Ethernet local area network (LAN) operating at 10, 100, or 1000 Mbps. Up to 10 features per z14 ZR1 server can be ordered.
0085	Trusted Key Entry (TKE) rack-mounted workstation: The rack-mounted version of the TKE, which needs a customer-provided, standard 19-inch rack. It includes a 1U TKE unit and a 1U console tray (screen, keyboard, and pointing device). When smart card readers are used, another customer provided tray is needed. Up to 10 features per z14 ZR1 server can be ordered.
0879	TKE 9.0 Licensed Internal Code (LIC): Included with the TKE tower workstation FC 0086 and the TKE rack-mounted workstation FC 0085 for z14 ZR1. Earlier versions of TKE features (FCs 0842, 0847, 0097, and 0098) can also be upgraded to TKE 9.0 LIC.
0895	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.

Feature code	Description
0892	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).

Note: You might need a TKE workstation that includes TKE Smart Card Reader while you run on CCA mode to satisfy certain industry security standard.

For more information about the Crypto Express6S feature and the corresponding crypto features, see the *IBM z14 Technical Guide*, SG24-8451.

10.1.2 Planning for Crypto Express6S configuration

Note: Support for Crypto Express6S coprocessors that are available on z14 ZR1 processors is included for z/OS V2.3, z/OS V2.2, and z/OS V2.1 in the Cryptographic Support for z/OS V2R1 – z/OS V2R3 (HCR77C1) web deliverable. It can be downloaded from the z/OS downloads website.

ICSF enhancements in z/OS V2.3 for the Crypto Express6S updates include support to use the new algorithm support and to extend existing support for asymmetric algorithms. This support also requires firmware or Microcode Change Level (MCL) updates to the TKE and z14 ZR1 processors, which are considered as co-requisites.

For more information about the latest MCL bundle requirements, see the Driver-27 Exception Letter that is available on IBM Resource Link (IBM ID authentication required).

The z14 ZR1 always operates in LPAR mode. The concept of *dedicated coprocessor* does not apply to the PCIe adapter. A PCIe adapter, whether configured as coprocessor or accelerator, is made available to logical partitions 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 in the partition.

The z14 ZR1 allows for up to 40 logical partitions to be active concurrently.

Each PCIe adapter on a Crypto Express6S feature supports 40 domains, whether it is configured as a Crypto Express6S coprocessor or a Crypto Express6S accelerator.

For availability reasons, the minimum configuration consists of two Crypto Express6S features so that every potential logical partition can have access to two cryptographic adapters.

More Crypto Express6S features might be needed to satisfy application performance and availability requirements. Consider the following points:

- For availability, spread assignment of multiple PCIe adapters of the same type (accelerator or coprocessor) to one logical partition across features in multiple I/O domains.
- The use of retained private keys on a PCIe adapter that is configured as a Crypto Express6S coprocessor creates an application single point of failure. This point of failure exists because RSA-retained private keys are not copied or backed up.

- An intrusion latch exists 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, the PCIe adapter can first be disabled from the TKE workstation before you remove the feature from the system. In that case, 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. The PCIe adapter then can be enabled from the TKE and normal operations can be resumed.

  Plan the definition of domain indexes and cryptographic coprocessor numbers in the Candidate list for each logical partition to prepare the cryptographic configuration. You can also define or change that cryptographic definition dynamically to an active logical partition with a running system. For more information, see “Changing LPAR Cryptographic Controls function” on page 225.
- Crypto Express6S features can be installed concurrently when all physical requirements are fulfilled. Dynamically enabling a new PCIe adapter to a partition requires the following configurations:
  - At least one usage domain index is defined to the logical partition
  - The cryptographic coprocessor numbers be defined in the partition Candidate list
- The same usage domain index can be defined more than once across multiple logical partitions. However, the cryptographic coprocessor number that is coupled with the specified usage domain index must be unique across all *active* logical partitions.

  The same cryptographic coprocessor number and usage domain index combination can be defined for more than one logical partition. This feature can be used, for example, to define a configuration for backup situations. In this case, only one of the logical partitions can be active at any one time.
- Newly installed Crypto Express6S features are assigned coprocessor numbers sequentially during the power-on-reset that follows the installation.

  However, when a Crypto Express6S 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 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. Each row identifies a PCIe adapter, and each column identifies a domain index number. Each cell entry indicates the logical partition to be assigned the cryptographic coprocessor number that is coupled with the usage domain index.

Table 10-2 Planning for logical partitions, domains, and PCIe adapter numbers

	Domain index 0	**Domain index 1**	**Domain index 2**	**.../...**	**Domain index 39**
PCIe adapter 0	LP00 LP02	LP04	LP05		
PCIe adapter 1	LP01 LP02				
PCIe adapter 2					
…					
…					

	Domain index 0	Domain index 1	Domain index 2	.../...	Domain index 39
PCIe adapter 13					
PCIe adapter 14					
PCIe adapter 15					

Table 10-2 on page 219 lists the following characteristics:

- Logical partitions LP00 and LP01 use domain 0 (zero), but are assigned different PCIe adapters. No conflict exists and they can be concurrently active.
- Logical partition 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.
- Logical partitions LP04 and LP05 use different domain numbers for PCIe cryptographic adapter 0 (zero); therefore, no conflict exists. 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 logical partition. The combination of cryptographic coprocessor number and usage domain index must be unique across all *active* logical partitions.

For more information about the Crypto Express6S feature for IBM Z, see *IBM z14 Technical Guide*, SG24-8451.

10.1.3 Configuring Crypto Express6S

This section describes the steps for configuring Crypto Express6S for the IBM z14 ZR1 server.

The z14 ZR1 operates in LPAR mode only. For each logical partition that requires access to a PCIe adapter (configured as an accelerator or coprocessor), the required information must be defined in the partition Image profile. This technique ensures the correct use of the cryptographic features when the associated partition is activated.

Concurrent changes to the Crypto Express6S features and controls when the partition is already activated are provided by special functions at the Support Element (SE).

Checking whether CPACF DES/TDES enablement feature is installed

The z14 ZR1 FC 3863 enables the DES and TDES algorithms on the CPACF. It is one of the prerequisites for using the Crypto Express6S feature. You must verify whether the CPACF feature is properly 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 HMC.
2. Open the **System details** menu of the CPC at the SE workplace. The system details window opens (see Figure 10-1 on page 221).

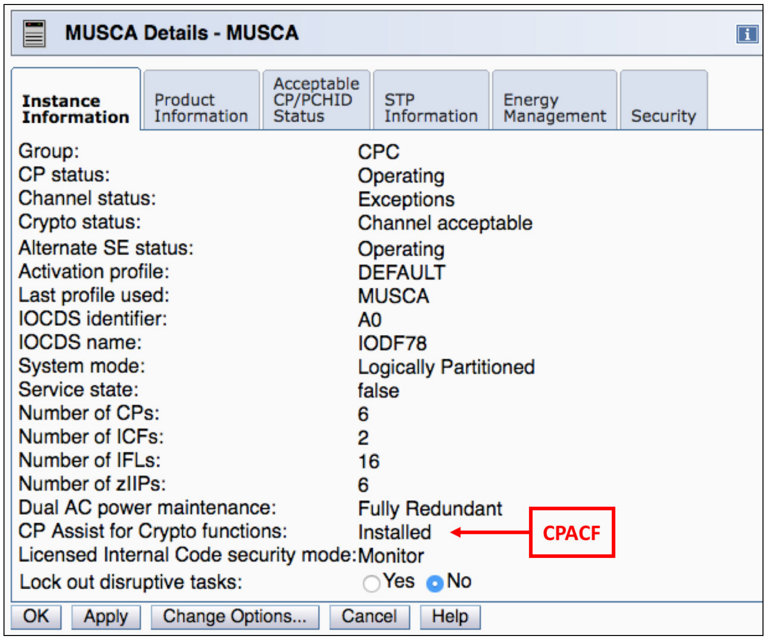

Figure 10-1 System details: CPACF installed

3. Click the **Instance Information** tab and verify that the CPACF DES/TDES enablement feature code 3863 is installed:
   - If the window shows the message `CP Assist for Crypto Functions: Installed`, the CPACF enablement FC 3863 is enabled.
   - If the window shows the message `CP Assist for Crypto Functions: Not installed`, FC 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 partition:

- Usage domain index
- Control domain index
- PCI Cryptographic Coprocessor Candidate List
- PCI Cryptographic Coprocessor Online List

These resources are defined 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 on page 222. After this definition is modified, any change to the image profile requires a DEACTIVATE and ACTIVATE of the logical partition for the change to take effect. Therefore, this kind of cryptographic definition is disruptive to a running system.

> **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. For more information about using this function, see "Changing LPAR Cryptographic Controls function" on page 225.

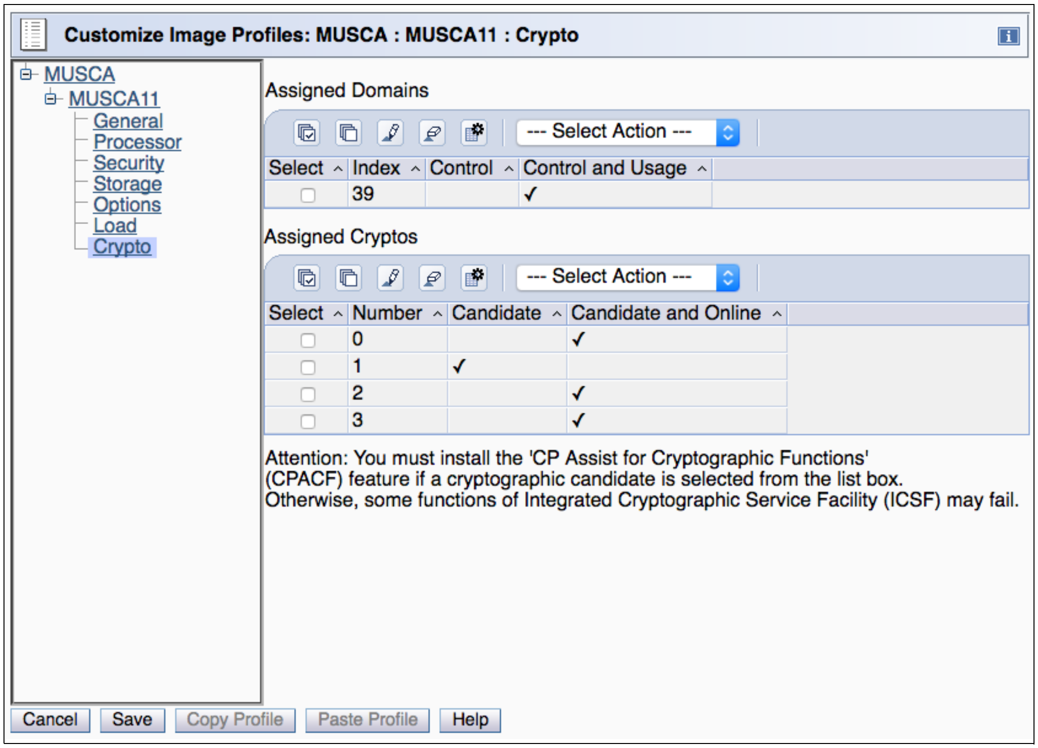

Figure 10-2 Customize Image Profiles: Crypto

The cryptographic resource definitions include the following meanings:

- Control Domain

  Identifies the cryptographic coprocessor domains that can be administered from this logical partition 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 logical partition 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 logical partition. From the list, select the coprocessor numbers (in the range 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 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 in use by another active logical partition, the activation of the logical partition fails (see Figure 10-3).

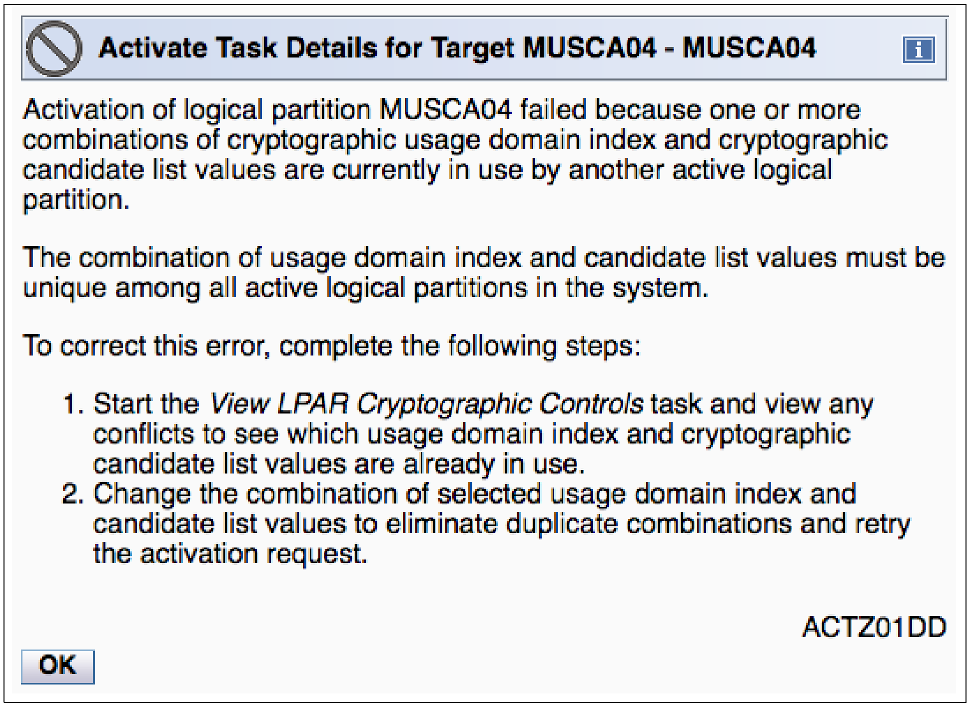

Figure 10-3 Activation of LPAR failed: ACTZ01DD

In this conflicting case, you must review the cryptographic information for all active logical partitions from the Summary tab of the View LPAR Cryptographic Controls task (see Figure 10-5 on page 225). Resolve the error based on the collected data by assigning a unique combination of PCIe adapter number and usage domain index number.

- Cryptographic Online list

  Identifies the cryptographic coprocessor numbers that are automatically brought online during logical partition 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 be configured online later to the partition by selecting **Configure On/Off** from the SE. For more information, see "Configuring a Crypto Express6S online or offline on a logical partition" on page 235.

  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 was selected in the partition's online list was 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 logical partition.

Cryptographic configuration using the Support Element

You can complete the following tasks from the SE:

- Display the PCI Cryptographic Configuration.
- Display the LPAR cryptographic controls (domain index and candidate or online lists for currently activate partitions).
- Reconfigure the coprocessor from or to the accelerator.
- Configure a cryptographic coprocessor and accelerator on or off to a logical partition.
- Change LPAR cryptographic controls to a logical partition.

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). The following options are available:

- View installed cryptographic features, with status and assigned PCHID and coprocessor numbers. Each PCIe adapter is assigned to a coprocessor number, in the range 0 - 15, 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.
- Start the release of coprocessor numbers. Remove the relationship only when a Crypto Express6S 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, which makes 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, the coprocessor number assignment remains.

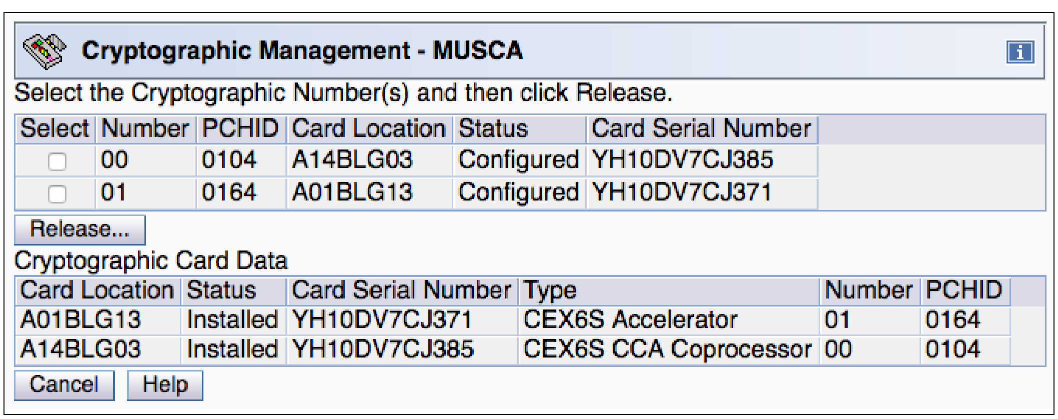

Figure 10-4 SE Cryptographic Management

Viewing LPAR Cryptographic Controls task

You can view active partition cryptographic definitions from the SE.

Select the CPCs and click **View LPAR Cryptographic Controls** in the Operational Customization window.

The resulting window displays the definition of Usage and Control domain indexes, and PCI Cryptographic candidate and online lists. The information is provided for active logical partitions only.

> **Tip:** You can review the PCI Cryptographic candidate lists and usage domain indexes that are assigned for all active logical partition from the Summary tab (see Figure 10-5). The usage domain index, in combination 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 a logical partition.

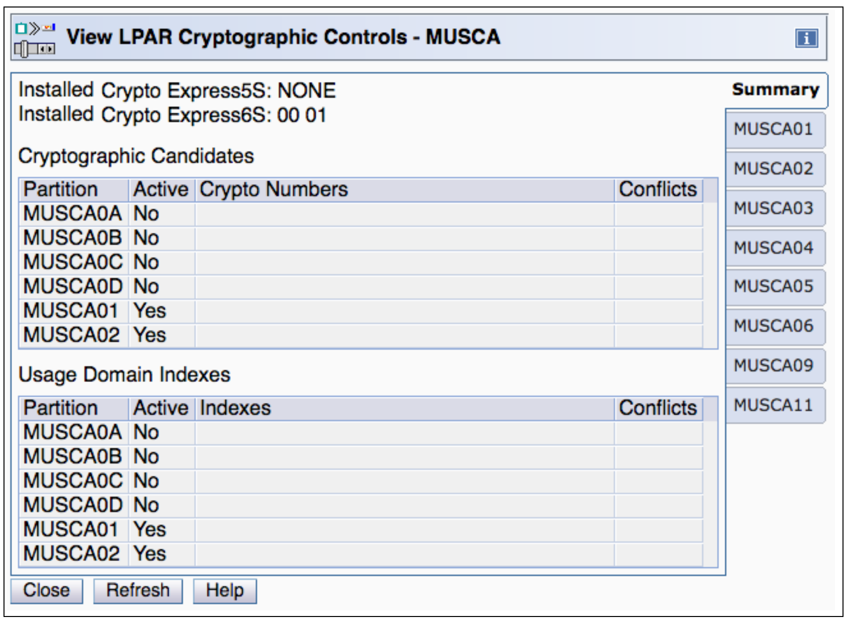

Figure 10-5 View LPAR Cryptographic Controls

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 use of the Configure On/Off task, which is described in “Configuring a Crypto Express6S online or offline on a logical partition” on page 235.

Changing LPAR Cryptographic Controls function

For each logical partition, you can define the following attributes:

- Usage domain index
- Control domain index
- Cryptographic Coprocessor Candidate list
- Cryptographic Coprocessor Online list

You can complete the following tasks by using the Change LPAR Cryptographic Controls function, which is included in the SE for the z14 ZR1:

- Add a cryptographic coprocessor to a logical partition for the first time.

- Add a cryptographic coprocessor to a logical partition that uses a cryptographic coprocessor.
- Remove a cryptographic coprocessor from a logical partition.
- 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 window at the SE. For more information about defining functions in the Image Profile, see “Logical partition cryptographic definition” on page 221.

Select the Control and Usage for each domain and the cryptographic Candidate and Online for each Crypto (see Figure 10-6).

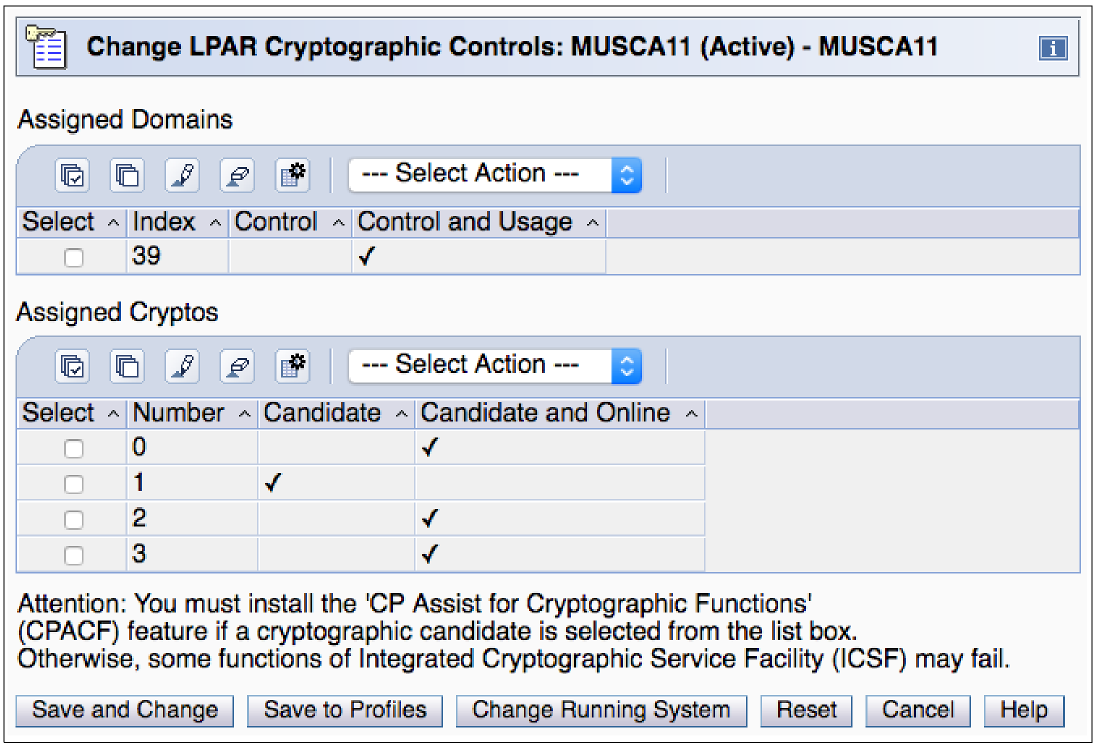

Figure 10-6 Change LPAR Cryptographic Controls: Change Running System

After selecting the appropriate options, you can complete the following 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 your changes are lost after a reactivation of the partition.

> **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 is displayed. After processing, a status window opens that indicates the result of a dynamic addition or change of a cryptographic definition to an LPAR (see Figure 10-7).

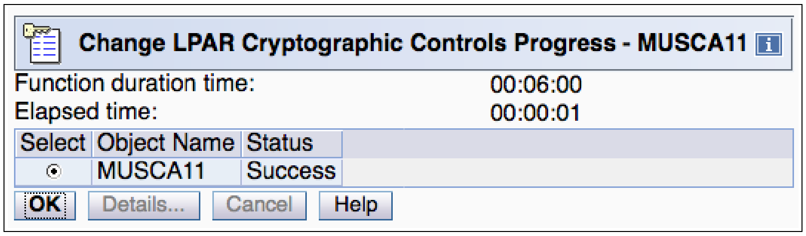

Figure 10-7 SE: Change LPAR Cryptographic Controls

Dynamic removal of the cryptographic definition

You can remove the cryptographic definition from a logical partition 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 you want to remove from a logical partition are offline (Standby). For more information about setting the cryptographic channel status, see “Configuring a Crypto Express6S online or offline on a logical partition” on page 235. 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.

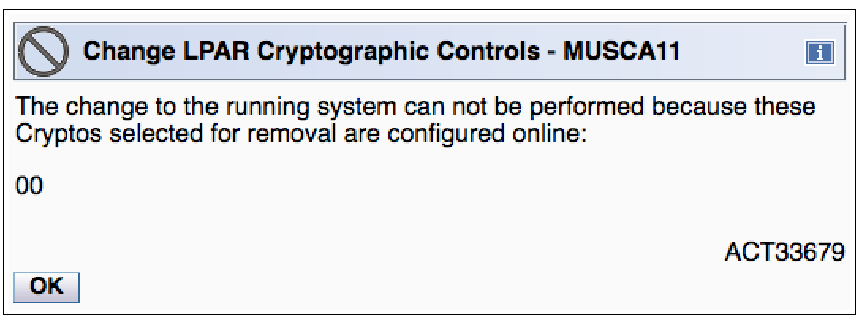

Figure 10-8 SE: Change LPAR Cryptographic Controls: ACT33679

   In addition to adding or changing cryptographic settings for a logical partition, you can remove the Control and Usage domains or Cryptographic Candidate lists for a logical partition from the Change LPAR Cryptographic Controls window (see Figure 10-6 on page 226).

   After clearing the definitions for a logical partition, 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 also cannot be used for the next image activation.

2. When you remove the only definition of the cryptographic lists, the Usage Domain Zeroize window opens (see Figure 10-9).

> **Consideration:** Because you cannot see all cryptographic information, including the usage domains for other logical partitions, you might need to check the information in **View LPAR Cryptographic Controls** window before you continue. For more information about zeroize, see “Reconfiguring the PCIe Adapter type” on page 228.

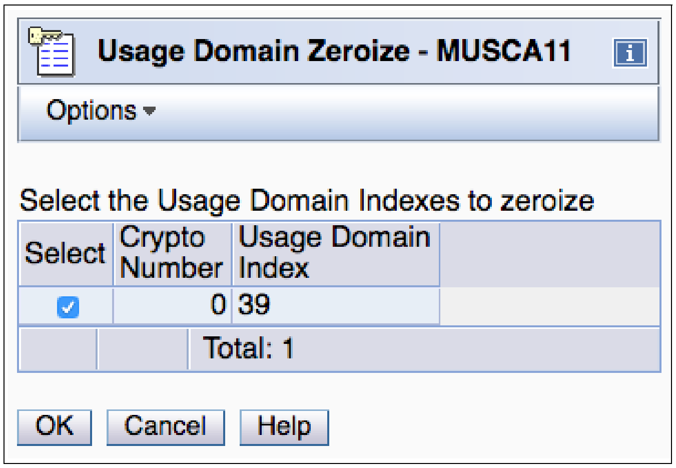

Figure 10-9 SE: Change LPAR Cryptographic Controls Zeroize

3. In the confirmation window (see Figure 10-10), click **OK** to dynamically change the cryptographic settings. After processing, a status window indicates the result of the dynamic change of cryptographic definition to an LPAR.

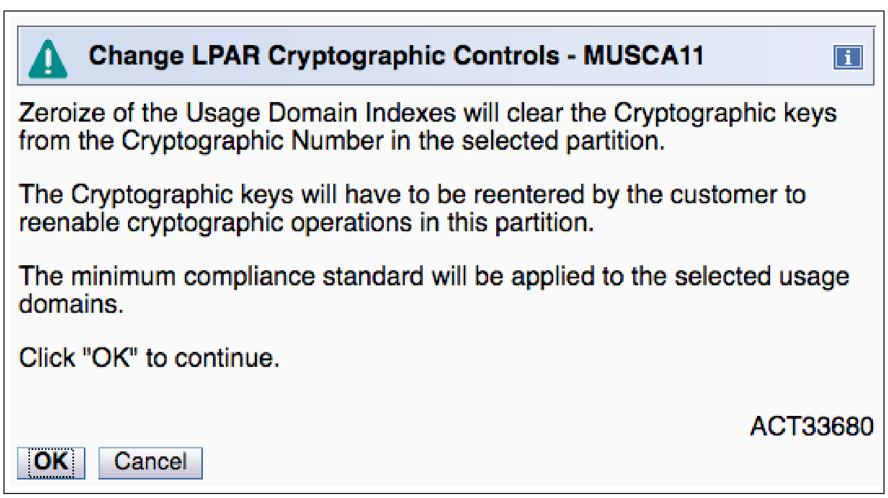

Figure 10-10 SE: Change LPAR Cryptographic Controls: ACT33680

Reconfiguring the PCIe Adapter type

Each PCIe Crypto Express6S feature can be configured as a coprocessor or an accelerator. Each Crypto Express6S feature can be set in the following configurations:

- Common Cryptographic Architecture (CCA) Coprocessor (CEX6C)
- Public Key Cryptography Standards #11 (PKCS) (EP11) Coprocessor (CEX6P)
- Accelerator (CEX6A)

Whether it is configured as a coprocessor or an accelerator, each PCIe Cryptographic adapter can be shared among 40 logical partitions.

Configuring a CCA coprocessor as an accelerator

During the installation of a Crypto Express6S feature, the PCIe Cryptographic adapter is configured by default as a CCA coprocessor. The reconfiguration is fully supported in Licensed Internal Code.

When a PCIe adapter is configured as a CCA coprocessor, it can still run accelerator functions, although much more slowly than when configured as 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 Express6S. For more information about configuring EP11 coprocessor, see “Configuring a CCA coprocessor as an EP11 coprocessor” on page 232.

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 Express6S online or offline on a logical partition” on page 235.
3. Select the number of the cryptographic coprocessor channel (see Figure 10-11) and click **Crypto Type Configuration**.

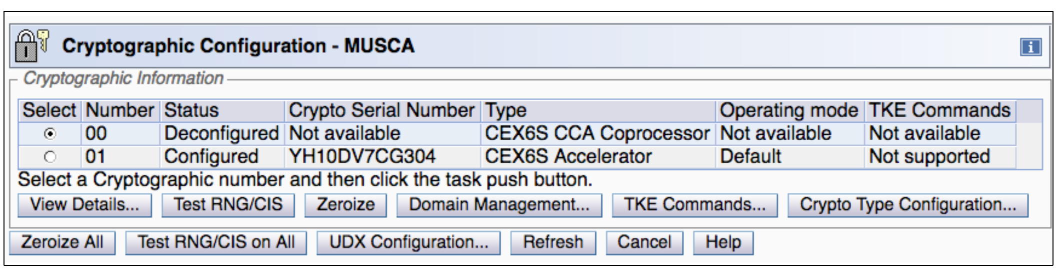

Figure 10-11 Cryptographic Configuration task (unconfigured)

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**.

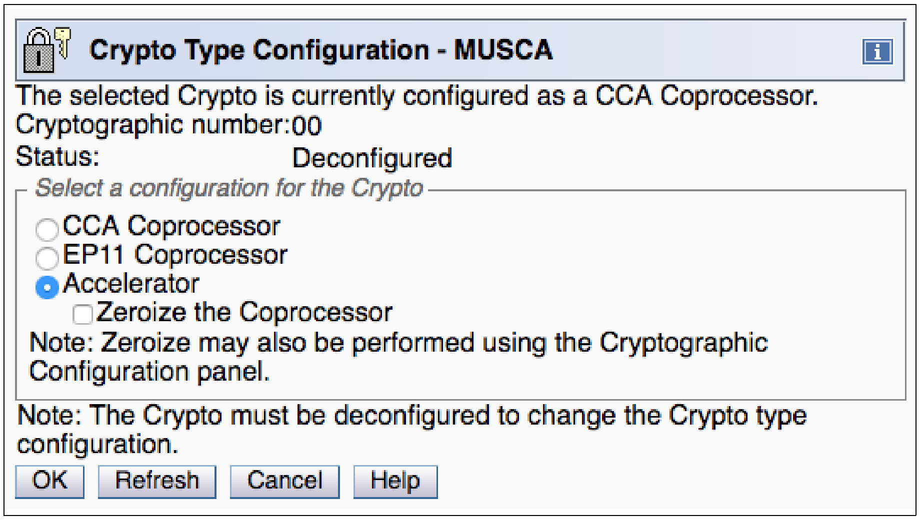

Figure 10-12 Crypto Type Configuration (CCA coprocessor to Accelerator)

By selecting Accelerator, you can zeroize the selected coprocessor by also selecting **Zeroize the Coprocessor** on the Crypto Type Configuration window. However, click to the **Zeroize the Coprocessor** option and then, click **OK**.

> **Important:** Zeroizing one or all cryptographic coprocessors clears their configuration data and all cryptographic keys. Zeroizing also erases configuration data from the SE hard disk drive (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 check box for each cryptographic channel.

5. Click **Yes** (see Figure 10-13).

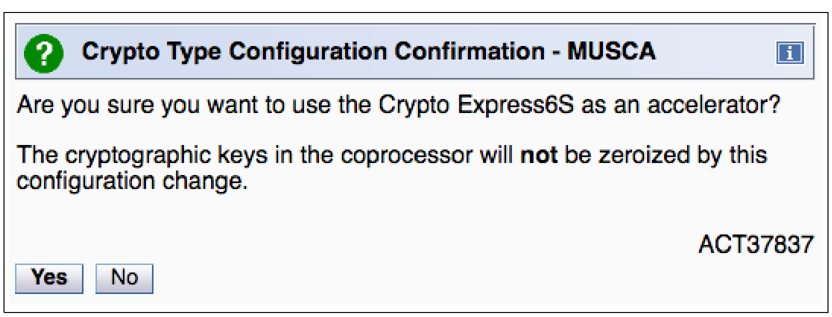

Figure 10-13 Crypto Type Configuration Confirmation for accelerator

6. Verify that your request 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 Express6S online or offline on a logical partition” on page 235.

After you perform this task and return to the Cryptographic Configuration window, the information that is shown in Figure 10-14 is displayed.

> **Note:** UDX support is *not* available for Crypto Express6S that is defined as an EP11 coprocessor and accelerator.

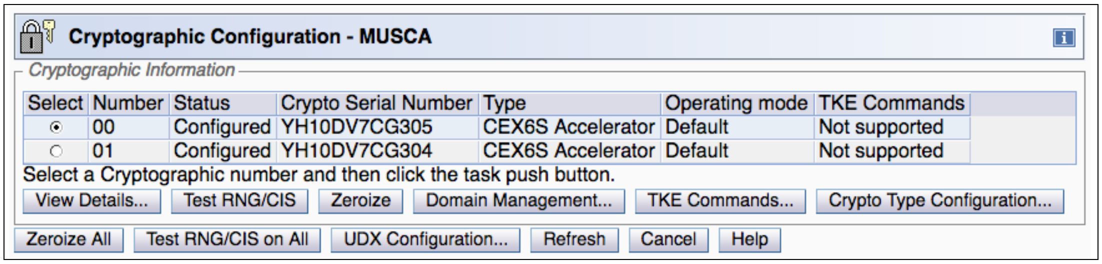

Figure 10-14 Cryptographic Configuration (Accelerator online)

8. Click **View Details** for detailed information (see Figure 10-15).

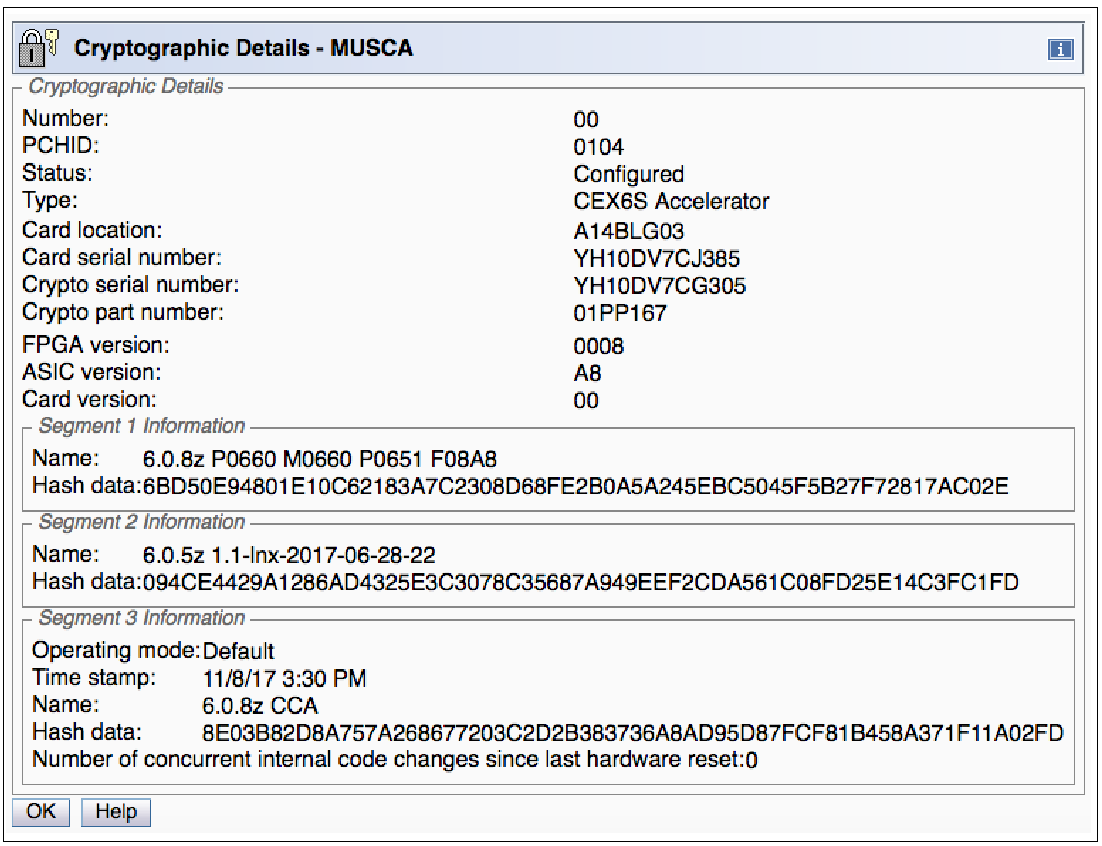

Figure 10-15 Cryptographic Details (Accelerator)

The Cryptographic Type is now a Crypto Express6S Accelerator. The adapter was *not* zeroized during the type-changing procedure.

The procedure for changing the type of the cryptographic configuration from an accelerator to a coprocessor is now 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 230.

The result of this change is shown in Figure 10-16.

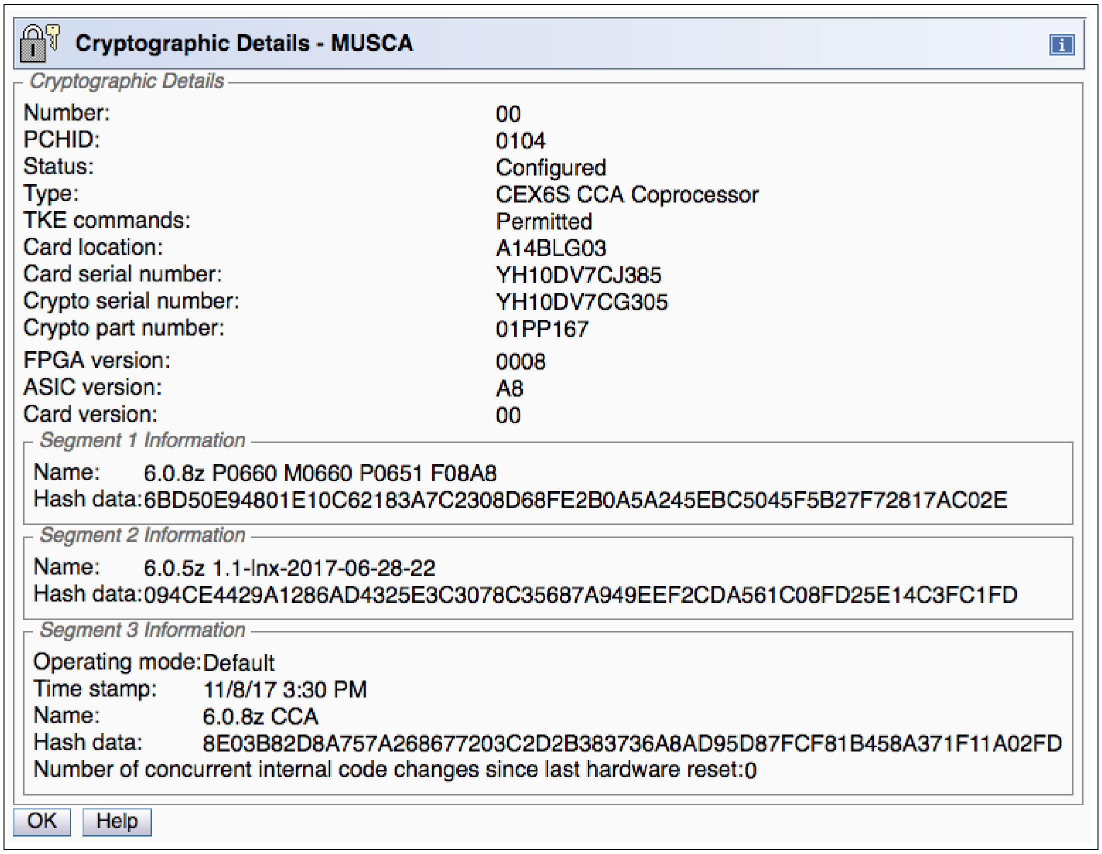

Figure 10-16 SE: Cryptographic Details (CCA Coprocessor)

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 includes cryptographic coprocessor adapters that you want to reconfigure and click **Cryptographic Configuration** in the CPC Configuration Group.
2. Because the reconfiguration is enabled only for PCIe adapters that are set to Off, be sure the PCIe Cryptographic adapter status for that cryptographic coprocessor channel is unconfigured (see Figure 10-12 on page 230).

   If necessary, set the PCIe Cryptographic adapter to Off for all partitions that are included in their candidate list. For more information about setting the PCIe Cryptographic adapter to Off, see “Configuring a Crypto Express6S online or offline on a logical partition” on page 235.
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 **OK**.

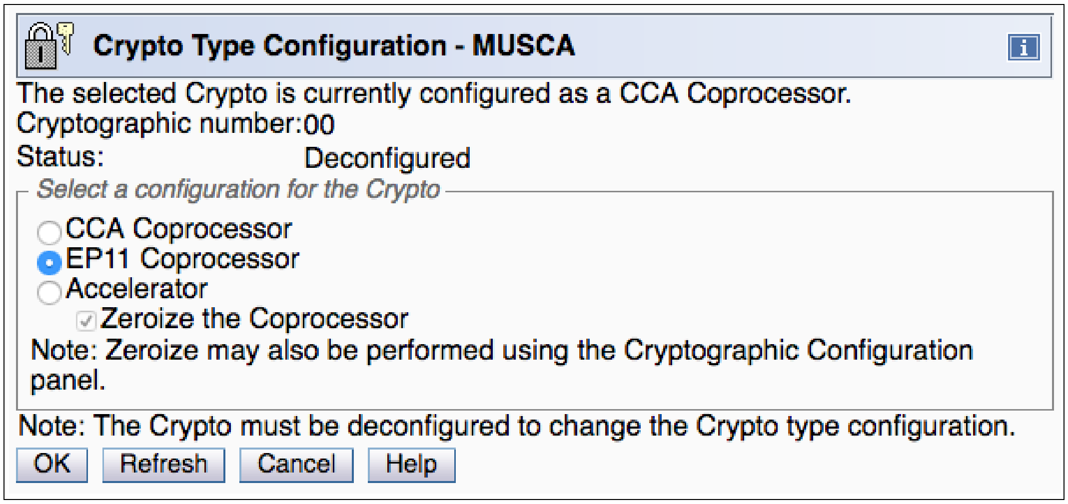

Figure 10-17 SE Crypto Type Configuration (CCA Coprocessor to EP11 Coprocessor)

5. Confirm your selection by clicking **Yes** (see Figure 10-18).

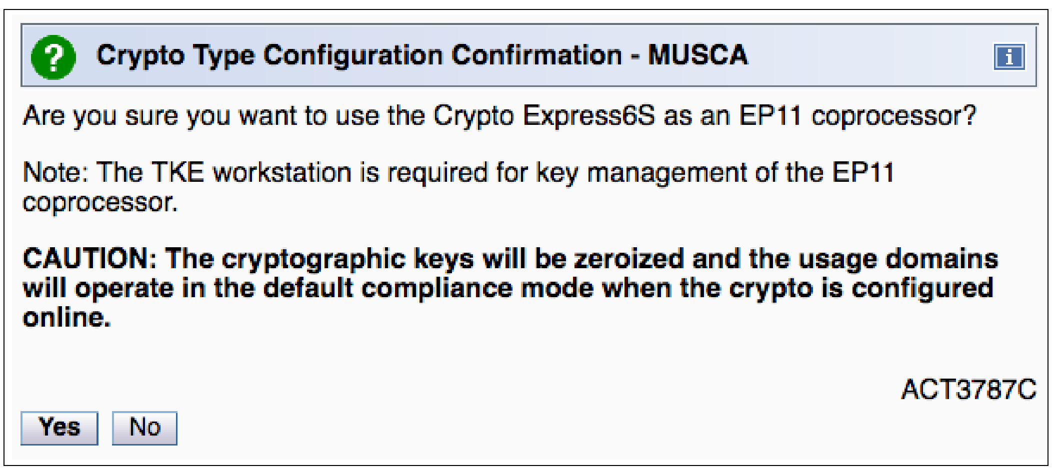

Figure 10-18 Crypto Type Configuration Confirmation for EP11 Coprocessor

6. Check that your request completed successfully. Click **OK** to return to the Crypto Type Configuration window.
7. Click **Cancel** in the Crypto Type Configuration window to return to the Cryptographic Configuration window. You can 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 Express6S online or offline on a logical partition” on page 235.

After you complete this task and return to the Cryptographic Configuration window, the information that is shown in Figure 10-19 is displayed.

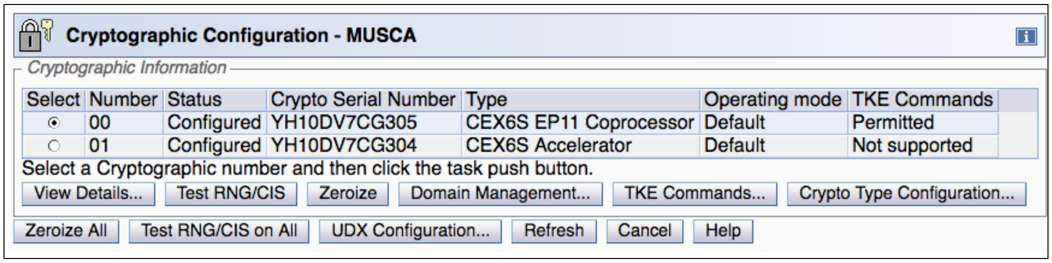

Figure 10-19 SE: Cryptographic Configuration (EP11 Coprocessor online)

8. Click **View Details** to display the detailed information, as shown in Figure 10-20.

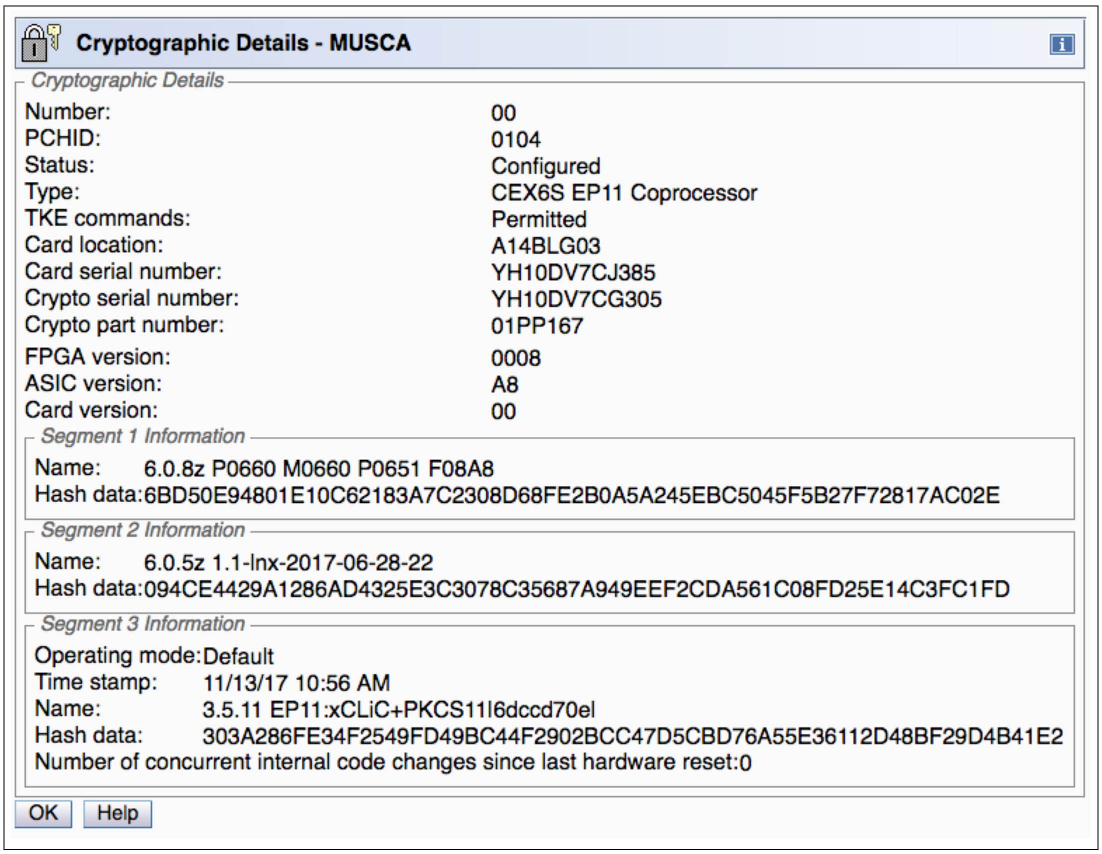

Figure 10-20 SE Cryptographic details (EP11 Coprocessor)

The Cryptographic Type is now a Crypto Express6S EP11 Coprocessor.

This process completes changing the type of the cryptographic configuration from a CCA Coprocessor to an EP11 coprocessor. To change the configuration back to 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 Express6S feature that is configured as an EP11 coprocessor, the TKE workstation is required.

Configuring a Crypto Express6S online or offline on a logical partition

For some changes to the cryptographic settings to the logical partition, you must configure the Crypto Express6S online or offline. This dynamic operation is not needed if you can reactivate (DEACTIVATE and ACTIVATE) the image for the logical partitions whose cryptographic online lists were updated.

Setting a Crypto Express6S to an online state

To set a Crypto Express6S online, complete the following steps:

1. From the SE, select the **System Management** function.
2. Select the server, click **Partitions**, and then, select the logical partition.
3. Click the **Cryptos** selection for the target logical partition.
4. In the contents of Cryptos page, select the Crypto IDs to be changed. Figure 10-21 shows that on server MUSCA to Logical Partition MUSCA11, two cryptographic coprocessors are defined: One CCA coprocessor (ID 00, physical Channel ID 0104), and one Accelerator (ID01, physical Channel ID 0164). The CCA coprocessor is online, and the Accelerator is offline. We now want to also set the Accelerator online.

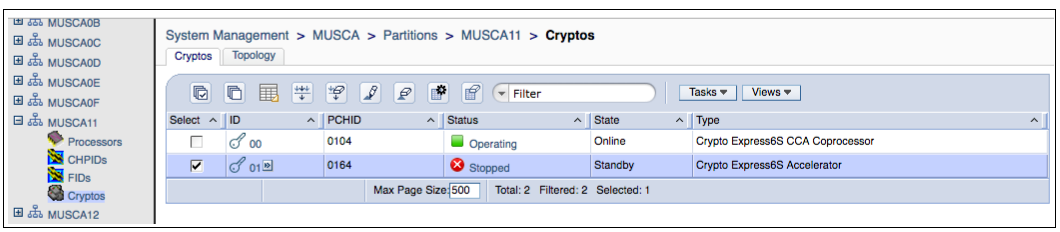

Figure 10-21 System Management: LPAR Crypto Selection, Standby

5. Select the cryptographic coprocessor and click **Tasks** → **Crypto Service Operations** → **Configure On/Off task** (see Figure 10-22). This task controls the online or offline (standby) state of a cryptographic processor for logical partitions that are defined in the cryptographic processor's candidate list.

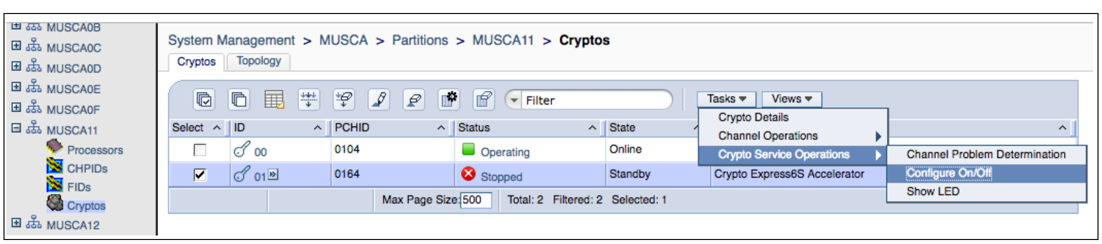

Figure 10-22 Crypto Service Operations: LPAR Crypto Selection, Configure Online

6. Select the cryptographic coprocessor channel number that you want, and then, click **Select Action** → **Toggle** to switch from `Standby` to `Online` (see Figure 10-23). If you want multiple cryptographic channels at the same time, select **Toggle All On**.

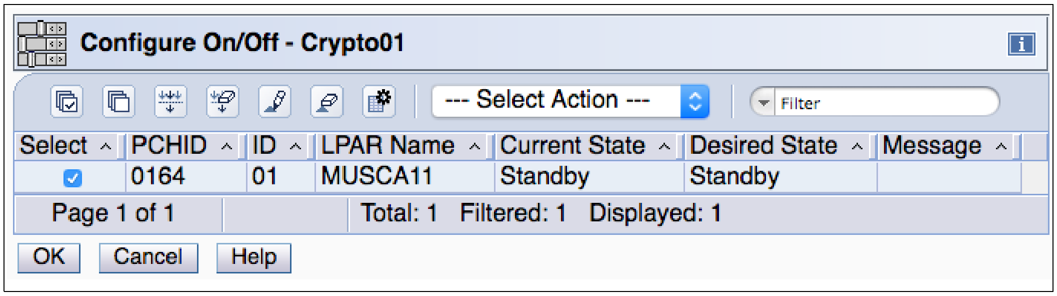

Figure 10-23 Config On/Off (Standby)

7. After confirming that your requested cryptographic coprocessor channel is set to the wanted state of `Online`, click **OK** (see Figure 10-24).

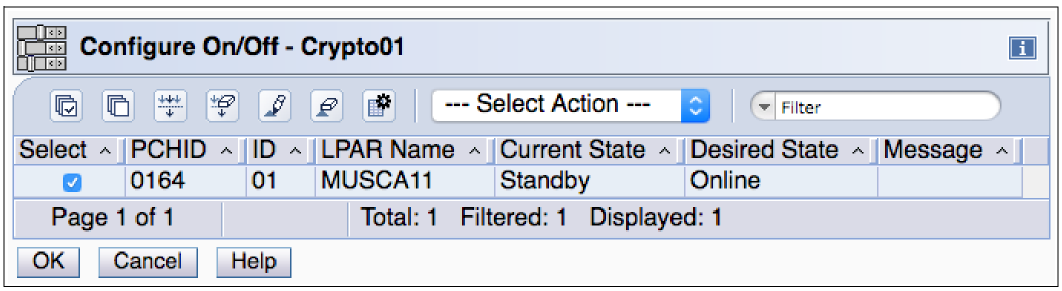

Figure 10-24 Configure On/Off (Standby to Online)

8. Confirm that your request is completed (see Figure 10-25). Click **OK**.

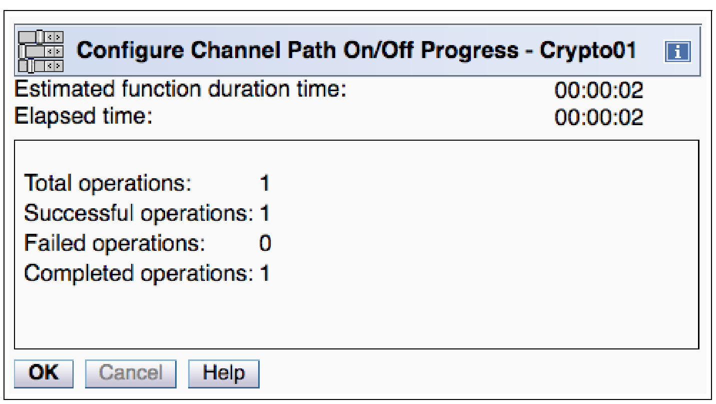

Figure 10-25 Configure On/Off (Standby to Online) completed

9. After you verify that the current state of the channels changed to `Online`, click **Cancel** to return.

You can view the contents of the Cryptos window of the logical partition to confirm that the cryptographic channels are now in the `Operating` status (see Figure 10-26).

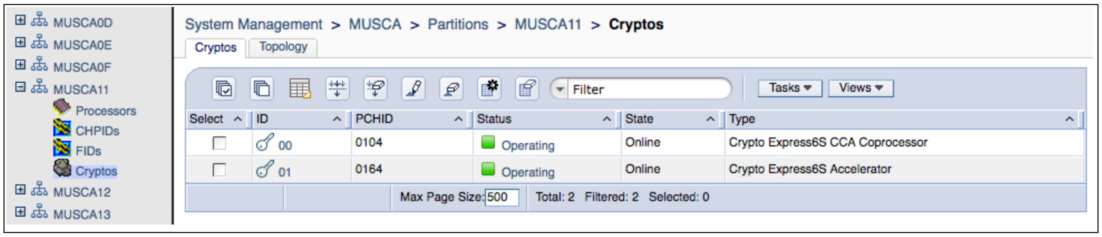

Figure 10-26 System Management: LPAR Crypto Selection, Online

Changing a cryptographic channel to standby (offline) status

To change the cryptographic channel status, complete the following steps:

1. Select the logical partition whose Crypto IDs you want to change to `Standby`. For example, select the Accelerator (01), that is in an online state. Click **Tasks** → **Crypto Service Operations** → **Configure On/Off task** (see Figure 10-27).

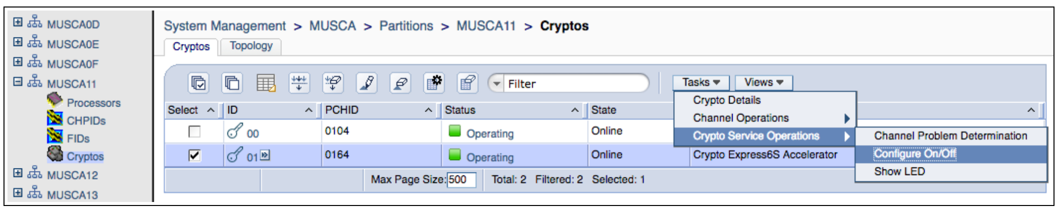

Figure 10-27 System Management: LPAR Crypto Selection, Configure Offline

2. Select the cryptographic coprocessor channel number that you want, and click **Select Action** → **Toggle All Standby** to switch from `Online` to `Standby` (see Figure 10-28).

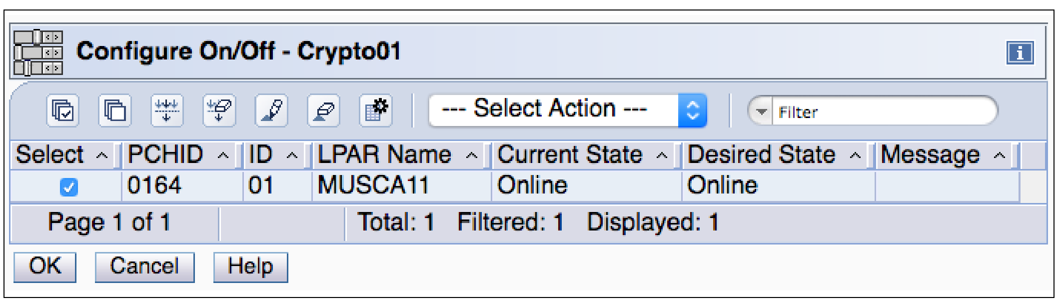

Figure 10-28 Config On/Off (Online)

3. After you confirm that the state for your requested cryptographic channel is `Standby`, click **OK** (see Figure 10-29).

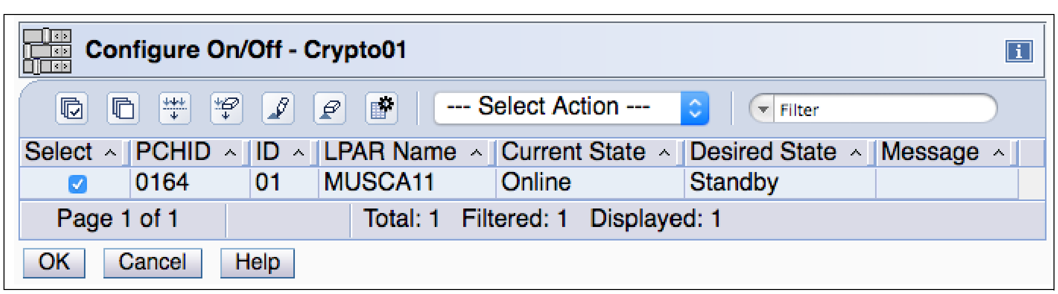

Figure 10-29 Configure On/Off (Online to Standby)

4. Because taking a cryptographic coprocessor offline can be disruptive to your application, a confirmation is required. The task must be confirmed by keying in the user password (Figure 10-30).

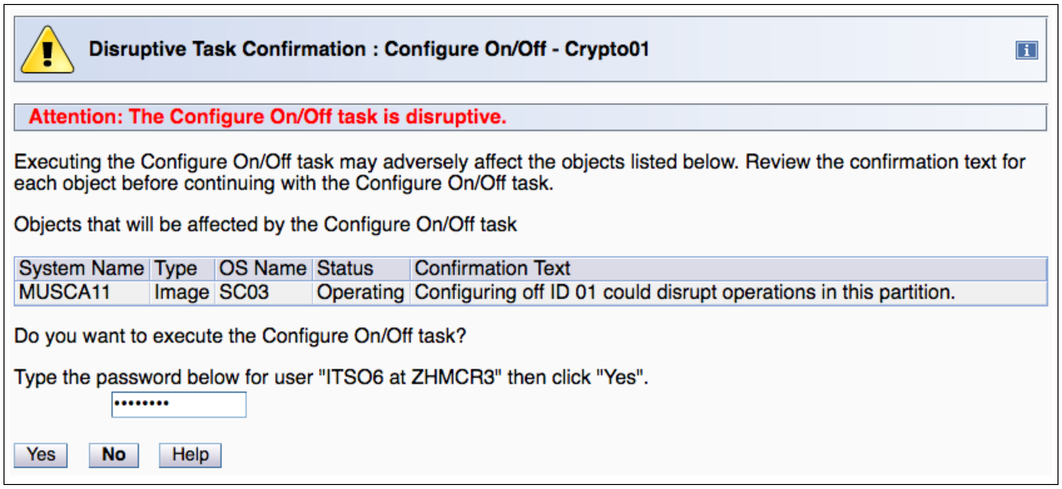

Figure 10-30 Configure On/Off (Online to Standby): Confirmation

5. Confirm that your request is completed (see Figure 10-31). Click **OK**.

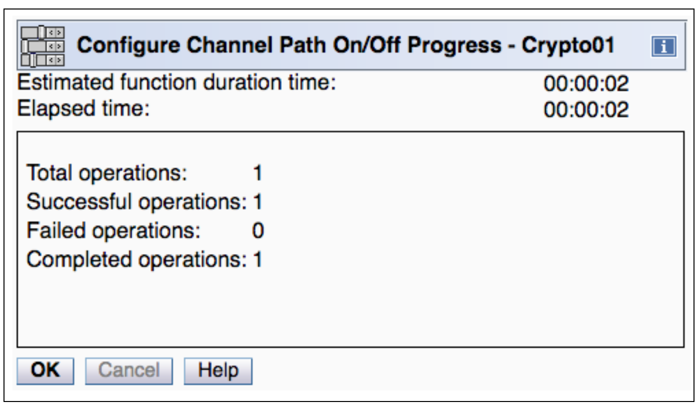

Figure 10-31 Config On/Off (Online to Standby) completed

10.1.4 Handling cryptographic coprocessors by using ICSF

Integrated cryptographic service facility (ICSF) provides an Interactive System Productivity Facility (ISPF) Coprocessor Management panel in which 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 z14 ZR1 SE.

It is not the purpose of this section to show how to create, load, and manage keys in the cryptographic adapters. For that information, see the ICSF literature. This section shows only how to activate and deactivate a cryptographic coprocessor, and how to display the hardware status.

From the ICSF main panel (see Figure 10-32), select option **1** to open the ICSF Coprocessor Management panel.

```
HCR77C1  -------------- Integrated Cryptographic Service Facility -------------
OPTION ===>
System Name:  SC03                           Crypto Domain: 39
Enter the number of the desired 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
     5650-ZOS Copyright IBM Corp. 1989, 2017.
     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 2 of 2
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
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   6P00    DV7CG305  Active                                         I
.   6A01    N/A       Active
****************************** Bottom of data *******************************
```

Figure 10-33 ICSF Coprocessor Management

When a coprocessor is configured offline to the logical partition 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 2 of 2
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
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   6P00    DV7CG305  Active                                         I
.   6A01    N/A       Offline
******************************* 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 the following 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.

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 2 of 2
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
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   6P00    DV7CG305  Active                                         I
d   6A01    N/A       Active
******************************* Bottom of data ********************************
```

Figure 10-35 ICSF Coprocessor Management (Online)

When a coprocessor is deactivated through IC0SF (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 2 of 2
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
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   6P00    DV7CG305  Active                                         I
.   6A01    N/A       Deactivated
******************************* Bottom of data ********************************
```

Figure 10-36 SF Coprocessor Management (Deactivated)

The Active/Deactivated status viewed from ICSF Coprocessor Management does not change the Online/Standby status that is set from the z14 ZR1 SE.

The ICSF Coprocessor Hardware Status panel is shown in Figure 10-37.

```
----------------- ICSF - PKCS #11 Coprocessor Hardware Status -----------------
COMMAND ===>                                                 SCROLL ===>
                                                            CRYPTO DOMAIN: 39

REGISTER STATUS                  COPROCESSOR 6P00

 Crypto Serial Number          : DV7CG305
 Status                        : ACTIVE
 Compliance Mode               : FIPS: 2009
                               : BSI:  NONE
 P11 Master Key
   New Master Key register     : EMPTY
     Verification pattern      :
                               :
   Current Master Key register : EMPTY
     Verification pattern      :
                               :

Press ENTER to refresh the hardware status display.
Press END   to exit to the previous menu.
```

Figure 10-37 ICSF Coprocessor Hardware Status

Help information from ICSF Coprocessor Management (see Figure 10-38 and Figure 10-39 on page 243) 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
------     ---------------------------------        -----------------------
   A        PCI Crypto Accelerator                   a, d
   X        PCI X Crypto Coprocessor                 a, d, e, k, r, s, v
  2C        Crypto Express2 Coprocessor              a, d, e, k, r, s, v
  2A        Crypto Express2 Accelerator              a, d,
  3C        Crypto Express3 Coprocessor              a, d, e, k, r, s, v
  3A        Crypto Express3 Accelerator              a, d
  4A        Crypto Express4 Accelerator              a, d
  4C        Crypto Express4 CCA coprocessor          a, d, e, k, r, s, v
  4P        Crypto Express4 PKCS #11 coprocessor     a, d, r, s
  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
  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

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' can not be combined with any other action characters.
The action character 'k' may be specified on only one coprocessor.
The action character 's' 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 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
  - Initializing stage 1: The feature has been detected by the configuration
                             task. No status is available.
```

Figure 10-38 Help for Coprocessor Management (part 1 of 2)

```
  - 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 take a 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.
  - 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 Key Store (CKDS, PKDS, or
        TKDS) and the master key is available for use
    C:  Master key Verification Pattern matches the Key Store, but the master
        key is not available for use
    E:  Master key Verification Pattern mismatch for Key Store or, for P11, no
        TKDS was specified in the options data set
    I:  The Master key Verification Pattern in the Key Store 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 zEnterprise Data Compression feature

This section describes the configuration of the zEnterprise Data Compression (zEDC) feature on an IBM z14 ZR1.

10.2.1 zEDC overview

The zEDC Express is an optional feature that is available with the zEC12, zBC12, z13, z13s, z14, and z14 ZR1. It is designed to provide hardware-based acceleration for data compression and decompression.

The zEDC Express feature is a native PCIe I/O card that can be installed in the PCIe I/O drawer, up to four zEDC Express features per drawer domain. A total of 1 - 8 features can be installed on the system. One PCIe adapter/compression coprocessor is available per feature, and a zEDC Express feature can be shared by up to 15 LPARs.

For more information about the zEDC feature, see *IBM z14 Technical Guide*, SG24-8451, and *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

10.2.2 Planning for zEDC configuration

For more information about planning, see 2.8.8, “Planning considerations for zEDC” on page 29.

10.2.3 Configuring zEDC

For more information about configuring a zEDC card by using HCD, see 15.2.4, “Defining a zEDC EXPRESS PCIe function” on page 379.

10.2.4 Handling zEDC

This section briefly describes how to manage the zEDC Express feature.

To verify that the IFAPRDxx member in your SYS1.PARMLIB concatenation is updated with the zEDC contents, use the **`D IQP`** command.

> **Note:** Schedule an IPL for every LPAR that you want zEDC active on. For more information, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

Example 10-1 shows the output of the **`DISPLAY IQP`** command that shows the status of the software feature.

Example 10-1 Display Feature Enablement: Enabled

```
DISPLAY IQP
IQP066I 15.09.38 DISPLAY IQP 961
zEDC Information
 MAXSEGMENTS:                4  (64M)
 Previous MAXSEGMENTS:     N/A
 Allocated segments:         1  (16M)
```

```
 Used segments:              0  (0M)
 DEFMINREQSIZE:              4K
 INFMINREQSIZE:             16K
 Feature Enablement:    Enabled
```

Verify that the hardware features are configured for the specific LPAR by displaying the status of the PCIe functions, as shown in Example 10-2.

Example 10-2 Display PCIe status - Hardware Accelerator: STNBY

```
DISPLAY PCIE
IQP022I 15.31.00 DISPLAY PCIE 981
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000A2  RoCE Express2            CNFG                   0100 0003 1
000000A3  RoCE Express2            CNFG                   0100 0004 2
00000102  8GB zHyperLink           CNFG                   0108 0003 1
00000103  8GB zHyperLink           CNFG                   0108 0004 2
000000D1  Hardware Accelerator     STNBY                  0140 0002
000000C1  Hardware Accelerator     STNBY                  0120 0002
```

The `STNBY` status denotes that the device is present and in standby mode and ready to be configured online.

Bring the device online by using the **`CONFIG PFID(xx),ONLINE`** command, as shown in Example 10-3.

Example 10-3 Configure PFID online

```
CONFIG PFID(D1),ONLINE
IEE504I PFID(D1),ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Verification of the PCIe status shows the new status of the PCIe functions, as shown in Example 10-4.

Example 10-4 Display PCIe status - Hardware Accelerator: ALLC

```
DISPLAY PCIE
IQP022I 15.45.52 DISPLAY PCIE 991
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000A2  RoCE Express2            CNFG                   0100 0003 1
000000A3  RoCE Express2            CNFG                   0100 0004 2
00000102  8GB zHyperLink           CNFG                   0108 0003 1
00000103  8GB zHyperLink           CNFG                   0108 0004 2
000000D1  Hardware Accelerator     ALLC    0012  FPGHWAM  0140 0002
000000C1  Hardware Accelerator     STNBY                  0120 0002
```

Notice that address space FPGHWAM (see Example 10-5) was started automatically by the system. Also, address space PCIe is started. Both address spaces are mandatory to handle PCIe functions.

Example 10-5 Display PCIe status extended format

```
DISPLAY PCIE,PFID=D1
IQP024I 07.47.23 DISPLAY PCIE 257
```

```
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000D1  Hardware Accelerator     ALLC    0012  FPGHWAM  0140 0002
 CLIENT ASIDS: NONE
 Application Description: zEDC Express
 Device State: Ready
 Adapter Info - Relid: 00000B  Arch Level: 03
                Build Date: 02/26/2014  Build Count: 00
 Application Info - Relid: 000000  Arch Level: 02
```

The **DISPLAY PCIE,PFID=D1** command that is shown in Example 10-5 displays the status of the zEDC Express feature as up and ready for use.

By design, when you must configure the feature to be offline, you must use the `FORCE` option because zEDC Express feature is always allocated by the `FPGHWAM` address space. Example 10-6 shows the configure offline command that uses the force option.

Example 10-6 Configure offline using force

```
CONFIG PFID(D1),OFFLINE,FORCE
IEE505I PFID(D1),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Displaying the PCIe status after the **CONFIG PFID(xx),OFFLINE,FORCE** command now shows the feature as STANDBY (see Example 10-7).

Example 10-7 Status after configure offline using force

```
DISPLAY PCIE
IQP022I 08.01.38 DISPLAY PCIE 284
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000A2  RoCE Express2            CNFG                   0100 0003 1
000000A3  RoCE Express2            CNFG                   0100 0004 2
00000102  8GB zHyperLink           CNFG                   0108 0003 1
00000103  8GB zHyperLink           CNFG                   0108 0004 2
000000D1  Hardware Accelerator     STNBY                  0140 0002
000000C1  Hardware Accelerator     STNBY                  0120 0002
```

10.3 Virtual Flash Memory

This section describes the configuration of the Virtual Flash Memory (VFM) feature on an IBM z14 ZR1.

10.3.1 VFM overview

The VFM is the replacement for the Flash Express features that were available on the zEC12, zBC12, z13, and z13s. VFM offers up to 2.0 TB of virtual flash memory in 512 GB increments. No application changes are required to change from Flash Express to VFM.

VFM is designed to help improve availability and handling of paging workload spikes when running z/OS V2.1, V2.2, or V2.3, or on z/OS V1.13. With this support, z/OS is designed to help 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 is also designed to help eliminate delays that can occur when collecting diagnostic data during failures. VFM can also be used in coupling facility images to provide extended capacity and availability for workloads that use IBM WebSphere MQ Shared Queues structures.

VFM can help organizations meet their most demanding service level agreements and compete more effectively. VFM is easy to configure, and provide rapid time to value.

For more information about the VFM feature, see the *IBM z14 Technical Guide*, SG24-8451.

10.3.2 Planning for VFM configuration

For planning considerations, see “Planning considerations for Virtual Flash Memory” on page 19.

10.3.3 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 to an activated LPAR. Therefore, the **Manage Flash Allocation** selection on the HMC is *not* supported for z14 ZR1.

Be aware of the following considerations when you allocate Virtual Flash Memory 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 allows more memory to be configured online, up to the maximum GB that is defined in this window, or up the maximum VFM 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 z14 ZR1.
- 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 click **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.

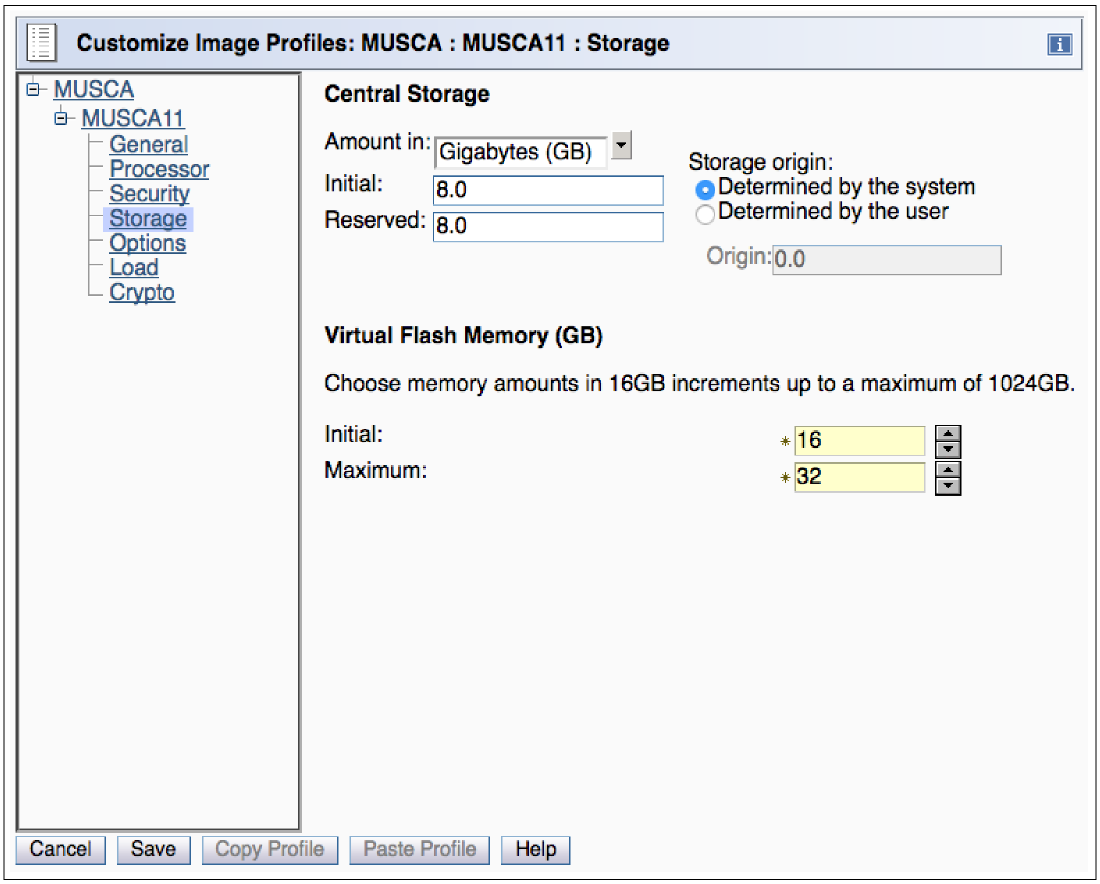

Figure 10-40 Virtual Flash Memory Allocation in Image Profile

In this example, the initial value is set to 16 GB, and the maximum value is set to 32 GB. The z14 ZR1 features two VFM features installed, which allows a maximum of 1024 GB allocated to the LPAR.

These definitions do not change the settings of a running LPAR. They are used only at the activation of the LPAR. A newly activated LPAR comes up with the specified amount of initial memory VFM.

If in the image activation profile of an LPAR an amount of initial VFM storage is specified that is greater than the available amount of deallocated VFM on the z14 ZR1, the activation of this LPAR fails with message ACTZ01EB, as shown in Figure 10-41.

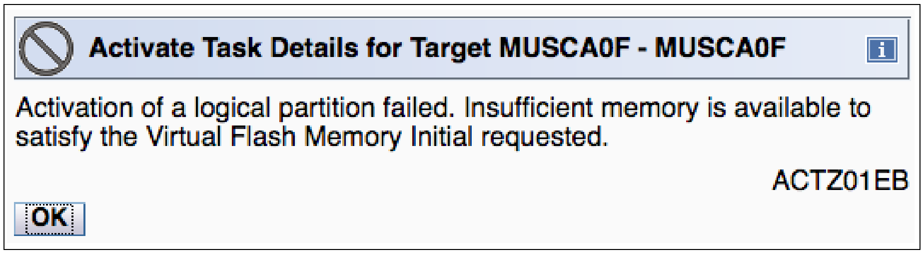

Figure 10-41 Insufficient VFM available, ACTZ01EB

> **Note:** For a Coupling Facility LPAR, it is also possible to define an initial value and a higher maximum value for VFM in the image profile. However, it does not make sense to set the maximum value higher than the initial value because the CFCC does not support any command to set any reserved memory online.

10.3.4 VFM management

The memory allocation of a z14 ZR1 is shown on the SE in the Storage Information window. To view it, click the server and select **Operational Customization** → **Storage Information**. Then, the Base System Storage™ Allocation window is displayed, as shown in Figure 10-42.

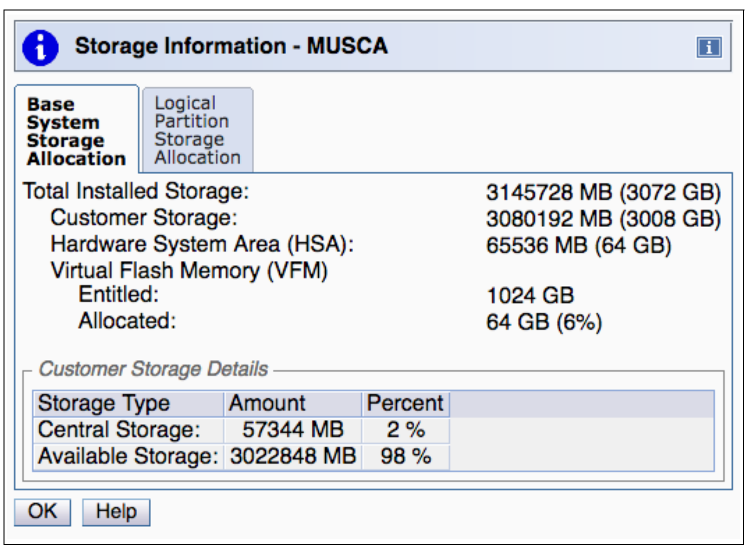

Figure 10-42 Storage Information: Base System Storage Allocation

The z14 ZR1 server in our examples (MUSCA) includes 1024 GB installed, of which 64 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 is listed (which were specified in the image activation profile), and also the currently allocated amount, as shown in Figure 10-43.

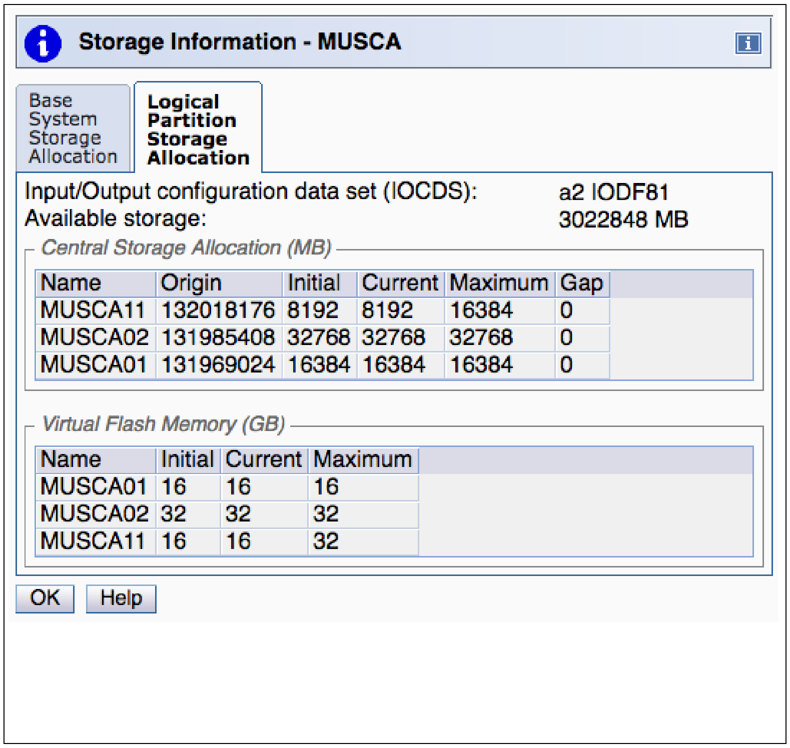

Figure 10-43 Storage Information: Logical Partition Storage Allocation

In z/OS for Flash Express and VFM, the PAGESCM parameter is supported in IEASYSxx. The syntax is shown in Example 10-8. This parameter determines whether and how much storage-class memory (SCM) is made available to an LPAR at IPL time.

Example 10-8 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. More SCM is allocated on an as-needed basis if use 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 SCM is 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 configure SCM online or offline to an LPAR (see Example 10-9).

Example 10-9 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 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 includes 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 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 the following 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 the following 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 amount of SCM 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 page data sets, which can lead to a loss of critical data during a DASD IBM HyperSwap® scenario.
> - SCM is used for paging large pages. If not enough SCM exists, 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 SCM or a list of ranges separated by commas identified by dddM|G|T-dddM|G|T; for example, 0G-16G, 32G-64G. The starting and ending addresses for each range of SCM must be multiples of the increment size.

The **`DISPLAY ASM`** and **`DISPLAY M`** commands include the following enhancements to display information and status that are related to Virtual Flash Memory:

`DISPLAY ASM`	Lists SCM status along with paging data set status.
`DISPLAY ASM,SCM`	Displays a summary of SCM usage.
`DISPLAY M=SCM`	Displays SCM online/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 the result of how ASM perceives its use of the cache of available SCM block IDs that ASM maintains. To ASM, some block IDs are not in use because they were not yet assigned to page out requests. However, to the **`DISPLAY M=SCM`** command processor, block IDs are in use because they were assigned to ASM for its use.

VFM storage can also be used by coupling facility LPARs running CFCC Level 22 on z14 ZR1, which is similar to Flash Express that be used by coupling facility LPARs running CFCC Level 19 (on zEC12 and zBC12) or CFCC Level 20 or 21 (on z13 an z13s). Systems without this support cannot connect to or rebuild a structure by using SCM storage.

In 10.3.3, “Configuring VFM” on page 247, we allocated an initial VFM of 16 GB to the LPAR MUSCA11, and a maximum VFM of 32 GB (see Figure 10-40 on page 248). Now, from MUSCA11 running the z/OS image SC03, we issue the IBM MVS™ **`DISPLAY IPLINFO,PAGESCM`** command. Example 10-10 shows the results.

Example 10-10 DISPLAY IPLINFO,PAGESCM

```
DISPLAY IPLINFO,PAGESCM
IEE255I SYSTEM PARAMETER 'PAGESCM': NOT_SPECIFIED
```

Because no `PAGESCM` parameter is 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), the initial amount of VFM that is specified is used automatically by z/OS for paging at IPL time. Likewise, if a specific amount is specified, this amount is made available for paging.

From SC03, run the enhanced **`DISPLAY ASM`** and **`DISPLAY M`** commands to display Virtual Flash Memory SCM-related information and status. The result for each command is shown in Example 10-11.

Example 10-11 Display commands

```
DISPPLAY ASM
IEE200I 12.04.41 DISPLAY ASM 799
TYPE     FULL STAT   DEV  DATASET NAME
PLPA      28%   OK  9A0B  PAGE.SC03.PLPA
COMMON     0%   OK  9A0B  PAGE.SC03.COMMON
LOCAL      0%   OK  9A36  PAGE.SC03.LOCAL01
LOCAL      0%   OK  9AB6  PAGE.SC03.LOCAL02
LOCAL      0%   OK  9B36  PAGE.SC03.LOCAL03
SCM        0%   OK   N/A   N/A
PAGEDEL COMMAND IS NOT ACTIVE

DISPLAY ASM,SCM
IEE207I 12.05.26 DISPLAY ASM 801
STATUS      FULL             SIZE            USED         IN-ERROR
IN-USE        0%         4,194,304          20,247               0

DISPLAY M=SCM
IEE174I 12.06.04 DISPLAY M 803
STORAGE-CLASS MEMORY STATUS
32G DEFINED
ONLINE
 0G-16G
16G OFFLINE-AVAILABLE
1% IN USE
SCM INCREMENT SIZE IS 16G
DISPLAY M=SCM(DETAIL)
IEE174I 12.06.53 DISPLAY M 805
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
32G DEFINED
ADDRESS  IN USE  STATUS
     0G      1%  ONLINE
ONLINE: 16G  OFFLINE-AVAILABLE: 16G  PENDING OFFLINE: 0G
1% IN USE
SCM INCREMENT SIZE IS 16G
```

From these commands, you see that 32 GB of VFM is defined, but only 16GB is online, while the other 16 GB are offline-available.

To vary another 16 GB VFM online to the example LPAR, issue the **CONFIG SCM(xxG),ONLINE** command, as shown in Example 10-12. The amount of VFM configured online must be specified according to the supported increment size. From these displays, the supported increment size is 16G.

Example 10-12 CONFIG SCM(16G),ONLINE

```
CONFIG SCM(16G),ONLINE
IEE195I SCM LOCATIONS 16G TO 32G ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Issue 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 (see Example 10-13).

Example 10-13 Post configuration displays

```
DISPLAY ASM
IEE200I 12.10.14 DISPLAY ASM 845
TYPE     FULL STAT   DEV  DATASET NAME
PLPA      28%   OK  9A0B  PAGE.SC03.PLPA
COMMON     0%   OK  9A0B  PAGE.SC03.COMMON
LOCAL      0%   OK  9A36  PAGE.SC03.LOCAL01
LOCAL      0%   OK  9AB6  PAGE.SC03.LOCAL02
LOCAL      0%   OK  9B36  PAGE.SC03.LOCAL03
SCM        0%   OK   N/A   N/A
PAGEDEL COMMAND IS NOT ACTIVE

DISPLAY ASM,SCM
IEE207I 12.10.41 DISPLAY ASM 847
STATUS      FULL              SIZE              USED         IN-ERROR
IN-USE        0%         8,388,608            20,247                0

DISPLAY M=SCM
IEE174I 12.08.47 DISPLAY M 843
STORAGE-CLASS MEMORY STATUS
32G DEFINED
ONLINE
 0G-32G
0G OFFLINE-AVAILABLE
0% IN USE
SCM INCREMENT SIZE IS 16G

DISPLAY M=SCM(DETAIL)
IEE174I 12.11.46 DISPLAY M 849
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
32G DEFINED
ADDRESS  IN USE  STATUS
     0G      1%  ONLINE
    16G      0%  ONLINE
ONLINE: 32G  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 249 and Figure 10-43 on page 250), this change in LPAR MUSCA11 is reflected.

As shown in Figure 10-44, the amount of allocated VFM went up to 80 GB (compared to the 64 GB that is shown in Figure 10-42 on page 249).

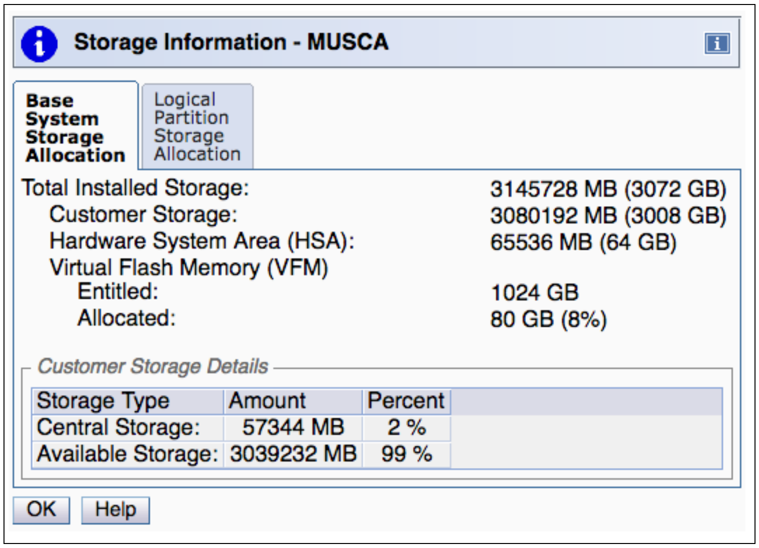

Figure 10-44 Base System Storage Allocation

In Figure 10-45, the amount of VFM that is allocated to LPAR MUSCA11 went up to 32 GB.

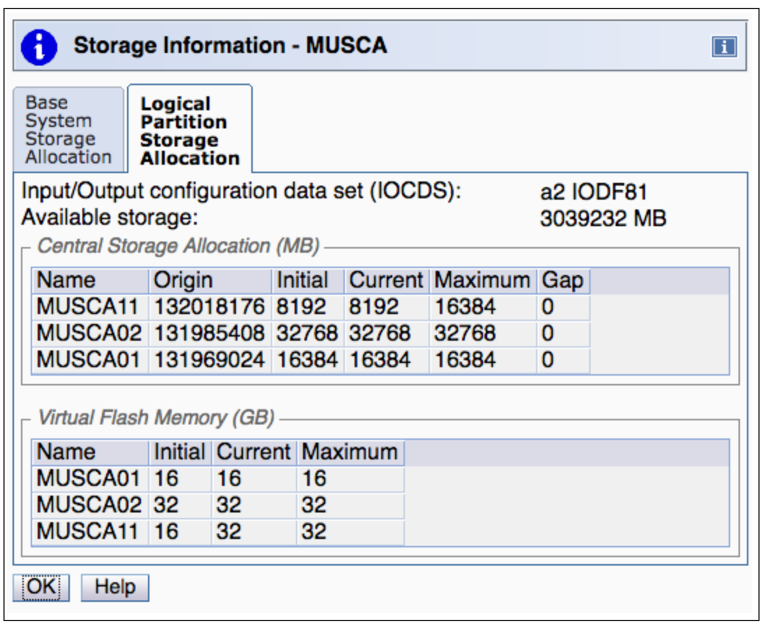

Figure 10-45 Logical Partition Storage Allocation

You also can set VFM offline, even to an amount that is lower than the initial value that is specified in the image activation profile. If for LPAR MUSCA11 the amount of online VFM is reduced to 0 GB by issuing **`CONFIG SCM(32G),OFFLINE`**, this process results in the Storage Information windows that are displayed in Figure 10-46 on page 256 and Figure 10-47 on page 256.

In Figure 10-46, the amount of allocated VFM was reduced to 48 GB.

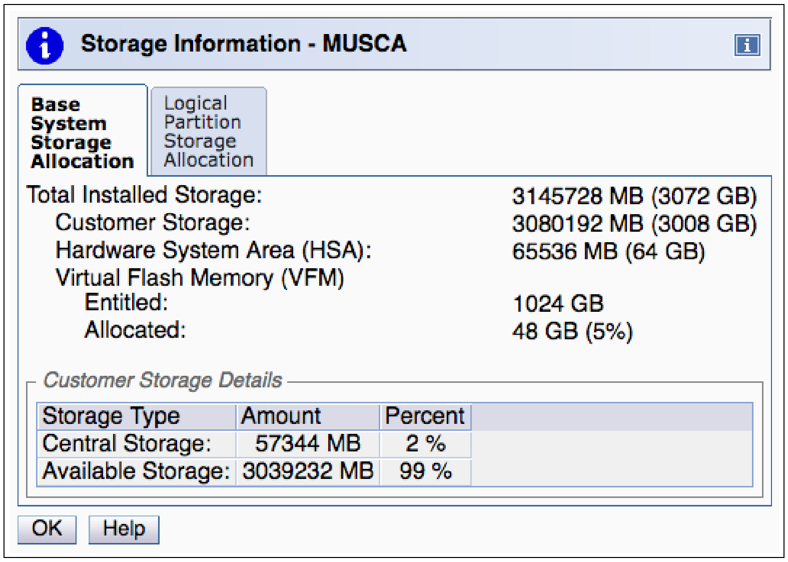

Figure 10-46 Base System Storage Allocation

In Figure 10-47, the amount of VFM that is allocated to LPAR MUSCA11 went down to 0 GB, which is lower than the initial 16 GB.

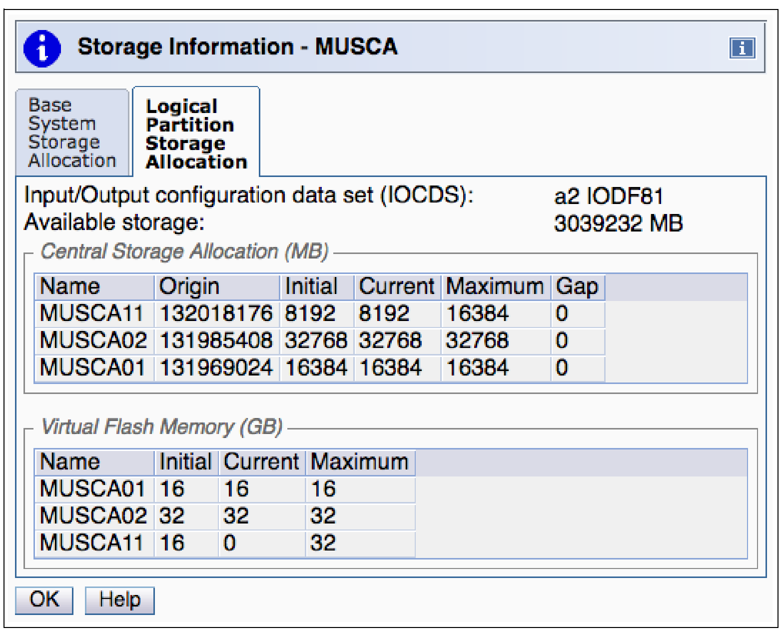

Figure 10-47 Logical Partition Storage Allocation

Note: An LPAR uses only the amount of VFM that is activated for that LPAR. VFM that is set offline by the operating system is returned to be used by other LPARs.

The allocation of VFM to a coupling facility LPAR is done in the same way as for z/OS LPARs and is described in 10.3.3, “Configuring VFM” on page 247. The amount of SCM that is allocated to a coupling facility LPAR can be displayed in the Operating System Messages window at the HMC.

For example, LPAR MUSCA1F, which allocated 32 GB of initial VFM storage, includes a message that shows the amount of SCM available, as shown in Example 10-14.

Example 10-14 CFCC messages with SCM

```
CF0280I CFCC Release 22.00, Service Level 00.30
        Built on 01/31/2018 at 15:50:00
        Code Load Features:
           Facility Operational Level: 22

CF0011I Coupling Facility is active with:
           2 CPs
           2 CF Receiver Channels
           0 CF Sender Channels
           1559 MB of allocatable storage
           32768 MB of Total SCM storage
```

The CF must know the algorithm of how the structure is used by the application. Currently, this algorithm is defined only for IBM MQ shared queues. To use this function, assign flash memory to your coupling facilities according to the procedure that is described next and update your structure definitions in your CFRM policy with the new parameter SCMMAXSIZE and SCMALGORITHM. For more information, see *z/OS MVS Setting Up a Sysplex*, SA23-1399.

IBM MQ for z/OS Version 7 or later allows the migration of IBM MQ shared queue objects to flash memory when structure usage exceeds the defined threshold. The IBM MQ objects are fetched back to real CF Storage when requested. This process provides an overflow capability for IBM MQ shared queues to handle workload peaks.

IBM RMF provides measurement data and reporting capabilities for VFM and Flash Express. The support enhances RMF Postprocessor and Monitor III reports with various new CF SCM statistics.

Coupling Facility SCM statistics are provided in the following reports:

- RMF Postprocessor Coupling Facility Activity (CF) report
- RMF Monitor III Coupling Facility Overview (CFOVER) report
- RMF Monitor III Coupling Facility Activity (CFACT) report

10.4 Shared Memory Communications over RDMA (SMC-R)

This section describes the configuration of the Shared Memory Communications over RDMA (SMC-R) that uses the 10GbE RoCE Express2 feature on an IBM z14 ZR1.

10.4.1 SMC-R overview

SMC-R uses the following IBM Z and industry standard communications technology:

- RDMA, which is based on queue pair (QP) technology that also uses an InfiniBand transport service type that is called reliable connected QPs (RC-QPs), which provide the following features:
  - Represent SMC Links in a logical point-to-point connection.

  - Transport data over unique RDMA network interface cards (RNICs) that are logically bound together to form Link Groups. Link Groups are used for high availability and load balancing needs.

  Ports in the IBM Z 10GbE RoCE Express2 feature (also referred to as RNICs) are used as the physical transport layer for RDMA.

- Single root I/O virtualization (SR-IOV) is a Peripheral Component Interconnect® Express (PCIe) standard that define extensions to PCIe specifications. SR-IOV enables sharing of 10GbE RoCE Express2 ports between LPARs in the z14 and z14 ZR1.

For more information about the 10GbE RoCE Express2 feature and SMC-R, see *IBM z14 Technical Guide*, SG24-8451, and *IBM z/OS V2R2 Communications Server TCP/IP Implementation Volume 1*, SG24-8360.

10.4.2 Planning for SMC-R configuration

For an overview of planning considerations, see “Shared Memory Communications - RDMA” on page 22.

10.4.3 Configuring SMC-R

The 10GbE RoCE Express and 10GbE RoCE Express2 features are native PCIe features; therefore, the following HCD and IOCP definition rules differ from a non-native PCIe card, such as OSA Express:

- PCIe Function Identifier (FID) must be defined in HCD or HCM to create IOCP input:
  - FID is a hexadecimal value (three heximal digits, range 000 - FFF), which specifies the PCIe function.
  - It *cannot* be assigned to a channel subsystem, so that any LPAR can be defined to a function.
  - It features a PARTITION parameter that dedicates it to one LPAR or allows reconfiguration among a group of LPARs. A function cannot be defined as *shared*.
  - In z/OS system commands, PCIe FID is represented as PFID.
- If the intended PCIe hardware supports multiple partitions, it has a decimal Virtual Function Identifier (VF=) in the range 1 - *n*, where *n* is the maximum number of partitions that the PCIe feature supports.
- Other parameters that are specific to the PCIe feature. For example, the 10GbE RoCE Express requires a Physical Network Identifier (PNETID=), and the new 10GbE RoCE Express2 feature supports a port identifier (PORT=).
- For function mapping to hardware, assign a Physical Channel Identifier (PCHID=) to identify the hardware feature in a specific PCIe I/O drawer and the slot to be used for the defined function. The following methods can be used:
  - Manually, by using the configurator (eCONFIG) PCHID report.
  - By using the CHPID Mapping tool and the eConfig Configuration Report File (CFR) input.

> **Note:** Unlike CHPIDs, multiple functions can be mapped to the same PCHID. This mapping is conceptually similar to mapping multiple InfiniBand coupling CHPIDs to the same adapter and port.

For more information about configuring an SMC-R connection by using HCD, see 15.2.3, “Defining a RoCE-2 PCIe function” on page 375.

10.4.4 SMC-R Management

This section introduces the z/OS commands that are related to the RoCE Express PCIe feature and shows the responses on our test system.

DISPLAY PCIE

You can use the **DISPLAY PCIE** command to display the following 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 10-15 shows an example of the **DISPLAY PCIE** command. You can confirm the FID and VFID that you defined. FID is represented as PFID (PCIE function identifiers).

Example 10-15 Example of D PCIE command

```
DISPLAY PCIE
IQP022I 13.32.12 DISPLAY PCIE 957
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000A2  RoCE Express2            CNFG                   0100 0003 1
000000A3  RoCE Express2            CNFG                   0100 0004 2
00000102  8GB zHyperLink           CNFG                   0108 0003 1
00000103  8GB zHyperLink           CNFG                   0108 0004 2
000000C1  Hardware Accelerator     ALLC    0012  FPGHWAM  0120 0002
000000D1  Hardware Accelerator     ALLC    0012  FPGHWAM  0140 0002
00000202  8GB zHyperLink           CNFG                   0168 0003 1
00000203  8GB zHyperLink           CNFG                   0168 0004 2
000000B2  RoCE Express2            CNFG                   0160 0003 1
000000B3  RoCE Express2            CNFG                   0160 0004 2
```

Example 10-16 shows an example of the **DISPLAY PCIE,PFID=pfid** command. After you define the new PCIe function, enter this command and confirm that its status is ACTIVE.

Example 10-16 Example of DISPLAY PCIE,PFID=pfid command

```
DISPLAY PCIE,PFID=0A3
IQP024I 13.36.48 DISPLAY PCIE 960
PCIE     0011 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
000000A3  RoCE Express2            CNFG                   0100 0004 2
 CLIENT ASIDS: NONE
 PNetID 1: PERFNET
```

Example 10-17 shows an example of the **DISPLAY PCIE,DD** command. You can confirm the details of the device drives that are installed in the system.

Example 10-17 Example of DISPLAY PCIE,DD command

```
DISPLAY PCIE,DD
IQP023I 13.42.11 DISPLAY PCIE 962
PCIE     0011 ACTIVE
DEV TYPE   DEVICE TYPE NAME
1014044B   Hardware Accelerator
10140613   8GB zHyperLink
15B36750   10GbE RoCE
15B31003   10GbE RoCE
15B31004   10GbE RoCE Express
15B31016   RoCE Express2
15B31014   40GbE RoCE Express2
101404ED   ISM
```

CONFIG command

You can use the **CONFIG** command to bring the PCIE function ID (PFID) online or offline.

Example 10-18 shows an example of **CONFIG PFID(xx),ONLINE** command.

Example 10-18 Example of CONFIG PFID(xx),ONLINE command

```
CONFIG PFID(A3),ONLINE
IEE504I PFID(A3),ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Example 10-19 shows an example of **CONFIG PFID(xx),OFFLINE** command.

Example 10-19 Example of CF PFID(x),OFFLINE command

```
CONFIG PFID(A3),OFFLINE
IEE505I PFID(A3),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

For more information about how to manage a RoCE Express feature, see *IBM z/OS V2R2 Communications Server TCP/IP Implementation Volume 1: Base Functions, Connectivity, and Routing*, SG24-8360.

10.5 Shared Memory Communications - Direct Memory Access

This section describes the configuration of the Shared Memory Communiction - Direct Memory Access (SMC-D) connections on an IBM z14 ZR1.

10.5.1 SMC-D overview

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

SMC-R requires a TCP/IP connection and preserves the entire network infrastructure. SMC-D is also a “hybrid” solution. It uses a TCP connection to establish the SMC-D connection. The TCP path can be through an OSA-Express port or HiperSockets connection.

A TCP option (called SMCD) controls switching from TCP to “out of band” SMC-D. The SMC-D information is exchanged within the TCP data stream. Socket application data is exchanged through ISM (write operations). The TCP connection remains established to control the SMC-D connection.

For more information about SMC-D, see *IBM z14 Technical Guide*, SG24-8451, and *IBM z/OS V2R2 Communications Server TCP/IP Implementation Volume 1*, SG24-8360.

10.5.2 Planning for SMC-D configuration

For more information about planning considerations for SMC-D, see “Shared Memory Communications - RDMA” on page 22.

10.5.3 Configuring SMC-D

For more information about configuring an SMC-D connection by using HCD, see 15.2.2, “Defining an ISM PCIe function” on page 372.

10.5.4 SMC-D management

From an operational standpoint, SMC-D is similar to SMC-R. However, SMC-D uses direct memory access (DMA) instead of an RDMA. It also uses a virtual PCI adapter that is called ISM rather than an RNIC. The ISM interfaces are associated with IP interfaces (for example, HiperSockets or OSA-Express) and are dynamically created, automatically started and stopped, and auto-discovered.

SMC-D over ISM does not use QP technology like SMC-R. Therefore, links and Link Groups that are based on QPs (or other hardware constructs) are not applicable to ISM. SMC-D protocol features a design concept of a “logical point-to-point connection” called an SMC-D link.

> **Note:** The SMC-D information in the `netstat` command displays is related to ISM link information (not Link Groups).

Internal Shared Memory technology

ISM is a virtual PCI network adapter that enables direct access to shared virtual memory, which provides highly optimized network communications for operating systems within the same IBM Z platform.

Virtual memory is managed by each z/OS (similar to SMC-R logically shared memory) following the existing IBM Z PCIe I/O translation architecture.

For more information about management of SMC-D, see *IBM z/OS V2R2 Communications Server TCP/IP Implementation Volume 1: Base Functions, Connectivity, and Routing*, SG24-8360.

10.6 IBM zHyperlink Express

This section describes the configuration of the zHyperlink Express feature on an IBM z14 ZR1.

10.6.1 IBM zHyperlink Express overview

IBM zHyperLink Express is a new, short distance, IBM Z I/O channel that is designed for up to 10x lower latency than High-Performance FICON. zHyperLink is intended to speed IBM Db2® for z/OS transaction processing and improve active log throughput. This feature is in the PCIe I/O drawer and is a two-port adapter used for short distance, direct connectivity between a z14 ZR1 and a DS8880. It uses PCIe Gen3 technology, with x16 lanes that are bifurcated into x8 lanes for storage connectivity.

The zHyperLink Express is designed to drive distances up to 150 meters (492.1 feet) and support a link data rate of 8 GigaBytes per second (GBps). A zHyperlink port is fully sharable between all partitions because 127 virtual functions/PFIDs per link are supported.

IBM zHyperLink dramatically reduces the latency of DASD I/Os by interconnecting the z14 ZR1 CPC directly to the I/O Bay of the DS8880. 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 Queueing Delays to resume it
- No host CPU cache disruption
- Reduced I/O service time

The zHyperLInk Express adapter takes one slot on z14 ZR1 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 z14 ZR1 server, thus resulting in up to 32 links.

FICON connectivity to each storage system is still required to be used for the following purposes:

- Initialization of the zHyperlink connection
- I/Os that are not eligible for zHyperlink
- Failback when a zHyperlink request fails (for example, cache miss or busy condition)

For more information about the zHyperlink feature, see *IBM z14 Technical Guide*, SG24-8451, and *IBM Z Connectivity Handbook*, SG24-5444.

10.6.2 Planning for zHyperlink Express configuration

For more information about planning considerations, see 2.8.4, “Defining the IBM zHyperLink Express” on page 21.

10.6.3 Configuring zHyperlink Express

Like the 10GbE RoCE Express and the 10GbE RoCE Express2 features, the zHyperlink Express card is a native PCIe I/O feature. Therefore, the HCD and IOCP definition rules for native PCIe cards also apply, as described in “Configuring SMC-R” on page 258.

zHyperLink ports are defined as PCI functions in the I/O configuration, where the PCHID represents the card and the port number represents the port. Each zHyperlink can be shared by up to 127 virtual functions. Generally, define four PFIDs per zHyperLink port per LPAR so that up to four simultaneous operations can be active on a link at a time (more operations cause link busy type conditions). Also, eight FICON CHPIDs can still be defined per Logical Control Unit (LCU) because zHyperLink does not reduce that number.

The association between zHyperLinks and storage systems is automatically discovered; therefore, no I/O configuration definition is required. The z14 ZR1 firmware discovers the storage subsystem during link initialization, and z/OS associates the zHyperlinks with the devices at IPL or vary online time.

For more information about configuring a zHyperling Express card by using HCD, see 15.2.5, “Defining a zHyperLink PCIe function” on page 383.

Managing zHyperlink Express

To enable IBM DB2® to use zHyperlinks, the DB2 zParm form zHyperlink must be enabled, as shown in Figure 10-48.

```
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 10-48 DB2 zParm: zHyperlink

The acceptable values for the DB2 zHyperlink Scope are:

ENABLE DB2 requests the zHyperlink protocol for all eligible I/O requests

DISBALE DB2 does not use the zHyperlink for any I/O requests

DATABASE DB2 requests the zHyperlink protocol for only database synchronous read I/Os

LOG DB2 requests the zHyperlink protocol for only log write I/Os

At GA time, zHyperlink is used by synchronous DB2 database reads only, which are provided with APAR PI77461 for DB2 V12.

To enable z/OS for zHyperlink read I/Os, the **ZHYPERLINK OPER=READ** statement must be added to the IECIOSxx parmlib member, as shown in Example 10-20.

Example 10-20 IECIOSxx parmlib enabled for zHyperlink read I/Os

```
VIEW      SYS1.PARMLIB(IECIOSFC) - 01.19                  Columns 00001 00072
Command ===>                                                 Scroll ===> PAGE
****** ***************************** Top of Data ******************************
000001 CTRACE(CTIIOS00)
000006 HYPERPAV=XPAV
000007 MIDAW=YES
000008 ZHPF=YES
000009 ZHYPERLINK OPER=READ
****** **************************** Bottom of Data ****************************
```

This process can also be done dynamically by entering the **SETIOS ZHYPERLINK,OPER=READ** console command. The corresponding display command is **DISPLAY IOS,ZHYPERLINK**, as shown in Example 10-21.

Example 10-21 DISPLAY IOS,ZHYPERLINK

```
DISPLAY IOS,ZHYPERLINK
IOS634I 14.20.06 IOS SYSTEM OPTION 998
ZHYPERLINK IS ENABLED FOR READ OPERATIONS
```

The **DISPLAY PCIE** command can be used to display the available PCIe function IDs for zHyperlink, as shown in Example 10-22.

Example 10-22 DISPLAY PCIE

```
DISPLAY PCIE
IQP022I 09.50.36 DISPLAY PCIE 291
PCIE     0010 ACTIVE
PFID      DEVICE TYPE NAME        STATUS  ASID  JOBNAME  CHID VFN  PN
00000001  Hardware Accelerator    ALLC    0011  FPGHWAM  0118 0002
00000011  Hardware Accelerator    ALLC    0011  FPGHWAM  0138 0002
00001304  8GB zHyperLink          ALLC    0017  IOSAS    01BC 0005 1
00001305  8GB zHyperLink          ALLC    0017  IOSAS    01BC 0006 1
00001306  8GB zHyperLink          ALLC    0017  IOSAS    01BC 0007 1
00001307  8GB zHyperLink          ALLC    0017  IOSAS    01BC 0008 1
00001384  8GB zHyperLink          CNFG                   01BC 0005 2
00001385  8GB zHyperLink          CNFG                   01BC 0006 2
00001386  8GB zHyperLink          CNFG                   01BC 0007 2
00001387  8GB zHyperLink          CNFG                   01BC 0008 2
00001004  8GB zHyperLink          ALLC    0017  IOSAS    013C 0005 1
```

Example 10-23 shows the **DISPLAY PCIE=pfid** command to display a specific zHyperlink PCIe function ID.

Example 10-23 DiSPLAY PCIE=pfid

```
DISPLAY PCIE,PFID=1304
IQP024I 09.54.25 DISPLAY PCIE 299
PCIE     0010 ACTIVE
PFID      DEVICE TYPE NAME         STATUS  ASID  JOBNAME  CHID VFN  PN
00001304  8GB zHyperLink           ALLC    0017  IOSAS    01BC 0005 1
 CLIENT ASIDS: NONE
 CU WWNN: 5005076306FFD680  CU Link Id: 0380
 S/W State: Allocated
 Port State: Operational
 CU Node Descriptor: 002107.981.IBM.75.0000000FAT71
```

The results of issuing the **DISPLAY M=CU(cun)** command against a control unit enabled for zHyperlink are shown in Example 10-24.

Example 10-24 DISPLAY M=CU(cun)

```
DISPLAY M=CU(9000)
IEE174I 09.56.31 DISPLAY M 307
CONTROL UNIT 9000
CHP                    8C   94   AC   D6   B4   BC   C4   CC
ENTRY LINK ADDRESS     C535 C543 C343 20F7 C907 C903 C713 C74A
DEST LINK ADDRESS      C330 C340 C530 C540 C730 C740 C930 C940
CHP PHYSICALLY ONLINE  Y    Y    Y    Y    Y    Y    Y    Y
PATH VALIDATED         Y    Y    Y    N    Y    N    Y    Y
MANAGED                N    N    N    N    N    N    N    N
ZHPF - CHPID           Y    Y    Y    Y    Y    Y    Y    Y
ZHPF - CU INTERFACE    Y    Y    Y    N    Y    N    Y    Y
INTERFACE ID           0010 0011 0012 .... 0140 .... 0142 0143
MAXIMUM MANAGED CHPID(S) ALLOWED = 0
DESTINATION CU LOGICAL ADDRESS = 00
CU ND             = 002107.981.IBM.75.0000000FAT71.0010
CU NED            = 002107.981.IBM.75.0000000FAT71.0000
TOKEN NED         = 002107.900.IBM.75.0000000FAT71.0000
WWNN              = 5005076306FFD680
FUNCTIONS ENABLED = ZHPF, ZHYPERLINK, XPAV
XPAV CU PEERS     = 9000, 9200
DEFINED DEVICES
  09000-0907F
DEFINED PAV ALIASES
  19000-1907F
ZHYPERLINKS
PFID      PCHID  Port  LinkId  S/W St   Port St
00001004  013C   01    0180    Alloc    Oper
00001005  013C   01    0180    Alloc    Oper
00001006  013C   01    0180    Alloc    Oper
00001007  013C   01    0180    Alloc    Oper
00001104  0178   01    0580    Alloc    Oper
00001105  0178   01    0580    Alloc    Oper
00001106  0178   01    0580    Alloc    Oper
00001107  0178   01    0580    Alloc    Oper
00001204  017C   01    0780    Alloc    Oper
```

```
00001205  017C   01    0780    Alloc    Oper
00001206  017C   01    0780    Alloc    Oper
00001207  017C   01    0780    Alloc    Oper
00001304  01BC   01    0380    Alloc    Oper
00001305  01BC   01    0380    Alloc    Oper
00001306  01BC   01    0380    Alloc    Oper
00001307  01BC   01    0380    Alloc    Oper
```

The results for the **DISPLAY M=DEV(devno)** command against a device enabled for zHyperlink are shown in Example 10-25.

Example 10-25 DISPLAY M=DEV(devno)

```
DISPLAY M=DEV(9000)
IEE174I 10.09.09 DISPLAY M 317
DEVICE 09000   STATUS=ONLINE
ENTRY LINK ADDRESS    C535 C543 C343 20F7 C907 C903 C713 C74A
DEST LINK ADDRESS     C330 C340 C530 C540 C730 C740 C930 C940
PATH ONLINE           Y    Y    Y    N    Y    N    Y    Y
CHP PHYSICALLY ONLINE Y    Y    Y    Y    Y    Y    Y    Y
PATH OPERATIONAL      Y    Y    Y    N    Y    N    Y    Y
MANAGED               N    N    N    N    N    N    N    N
CU NUMBER             9000 9000 9000 9000 9000 9000 9000 9000
INTERFACE ID          0010 0011 0012 .... 0140 .... 0142 0143
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 002107.981.IBM.75.0000000FAT71.0010
SCP TOKEN NED     = 002107.900.IBM.75.0000000FAT71.0000
SCP DEVICE NED    = 002107.900.IBM.75.0000000FAT71.0000
WWNN              = 5005076306FFD680
HYPERPAV ALIASES CONFIGURED = 128
ZHYPERLINKS AVAILABLE = 16
FUNCTIONS ENABLED = MIDAW, ZHPF, XPAV, ZHYPERLINK
```

This display command was enhanced with the new parameter **ZHYPERLINK** to show whether the device supports zHyperlink. The response is shown for a device that is capable for zHyperlink in Example 10-26.

Example 10-26 DISPLAY M=DEV(devno),ZHYPERLINK: Device enabled for zHyperlink

```
DISPLAY M=DEV(7000),ZHYPERLINK
IEE587I 14.57.37 DISPLAY M 356
DEVICE 07000   STATUS=ONLINE
DEVICE IS ENABLED FOR ZHYPERLINK
READ OPERATIONS ARE ENABLED
WRITE OPERATIONS ARE DISABLED FOR THE FOLLOWING REASON(S):
  ZHYPERLINK WRITES ARE DISABLED FOR THE SYSTEM
  CONTROL UNIT DOES NOT SUPPORT ZHYPERLINK WRITES
```

The result for a device that is not capable for zHyperlink is shown in Example 10-27.

Example 10-27 DISPLAY M=DEV(devno,ZHYPERLINK: Device not enabled for zHyperlink

```
DISPLAY M=DEV(265E),ZHYPERLINK
IEE587I 14.59.25 DISPLAY M 373
DEVICE 0265E   STATUS=ONLINE
```

```
DEVICE IS DISABLED FOR ZHYPERLINK FOR THE FOLLOWING REASON(S):
  CONTROL UNIT DOES NOT SUPPORT ZHYPERLINK
  THERE ARE NO ZHYPERLINKS AVAILABLE
WRITE OPERATIONS ARE DISABLED FOR THE FOLLOWING REASON(S):
  ZHYPERLINK WRITES ARE DISABLED FOR THE SYSTEM
```

11

Adding LPARs and OSCONFIGs

This chapter describes the steps for adding Logical Partitions and Operating System Configurations to your IODF. It also includes a list of potential configuration items and a short description on how to complete each item by using the Hardware Configuration Definition (HCD) tool.

Note: Not all of the configuration items that are described in this chapter are necessarily required for your installation. Also, the examples that are presented are not intended to be exhaustive.

For more information about how to use HCD and FICON, see the following publications:

- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804
- *FICON Planning and Implementation Guide*, SG24-6497

This chapter includes the following topics:

11.1 Defining more I/O by using HCD

When defining new I/O components in an IODF, certain definitions, such as OSCONFIGs, partitions, FICON Switches, Control Units, and Devices, must be made up front. After these items are defined, the following connections can be made:

- Defining another OSCONFIG
- Defining more operating system LPARs in a CSS

The I/O definitions that are described next use HCD to demonstrate the examples. The examples continue by using the example that was created in Chapter 5, “Production IODF and setting up the CPC” on page 85.

11.2 OSCONFIGs and Logical Partition definitions

This section describes defining Operating System Configurations, Logical Partitions, and FICON Switches.

11.2.1 Defining another OSCONFIG

The following prerequisites must be met for a new OSCONFIG:

- OSCONFIG name:
  - EDT ID (Eligible Device Table ID): Esoterics / VIO
  - Consoles
- To add devices to Esoterics and Consoles, they first must be added to an OSCONFIG.

To add an OSCONFIG by using the HCD tool, complete the following steps:

1. From the main HCD panel, select option **1.1. Operating system configurations**.
2. In the Command line, enter `add` (see Figure 11-1 on page 270) to add an OSCONFIG.
3. Complete the following updates and press Enter:
   - Update `OS configuration ID type` to ITSOTEST.
   - Update `Description` to ITSO test OSCONFIG for z14 ZR1.

```
+---------------- Add Operating System Configuration -----------------+
|                                                                     |
|                                                                     |
| Specify or revise the following values.                             |
|                                                                     |
| OS configuration ID  . . . . . ITSOTEST                             |
| Operating system type  . . . . MVS       +                          |
|                                                                     |
| Description  . . . . . . . . . ITSO test OSCONFIG for z14 ZR1_      |
|                                                                     |
| 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 Operating System Configuration: Add OSCONFIG

4. To add an EDT ID, enter `s` next to the new OSCONFIG (see Figure 11-2).

```
                   Operating System Configuration List             Row 1 of 34
Command ===> ___________________________________________ Scroll ===> CSR

Select one or more operating system configurations, then press Enter. To
add, use F11.

/ Config. ID  Type     Gen  Description                       D/R site OS ID
_ DBSV4SU4    MVS           z/OS DB Server 4
_ DBSV5SU4    MVS           z/OS DB Server 5
_ DBSV6SU4    MVS           z/OS DB Server 6
_ ITSO        MVS           All ITSO devices
s ITSOTEST    MVS           ITSO test OSCONFIG for z14 ZR1
_ PERF4SU4    MVS           z/OS Appl Server 4
```

Figure 11-2 Operating System Configuration: Add EDT

5. In the command line, enter `add` (see Figure 11-3) to add an EDT.

```
+-------------------------------- Add EDT --------------------------------+
|                                                                          |
|                                                                          |
| Specify the following values.                                            |
|                                                                          |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1         |
|                                                                          |
| EDT identifier . . . 00                                                  |
| Description  . . . . ________________________________                    |
|                                                                          |
|  F1=Help     F2=Split    F3=Exit     F5=Reset    F9=Swap    F12=Cancel   |
+--------------------------------------------------------------------------+
```

Figure 11-3 Operating System Configuration: Add EDT

6. Update `EDT identifier` to `00` and press Enter.
7. Enter `s (work with esoterics)` next to `EDT 00` and press Enter (see Figure 11-4).

```
+------------------------------- EDT List --------------------------------+
|   Goto  Backup  Query  Help                                              |
| ----------------------------------------------------------------------   |
|                                                             Row 1 of 1   |
| Command ===> _______________________________________ Scroll ===> CSR     |
|                                                                          |
| Select one or more EDTs, then press Enter. To add, use F11.              |
|                                                                          |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for z14             |
|                                                                          |
| / EDT Last Update By       Description                                   |
| s 00  2018-03-16  ITSO1    ________________________________              |
| *************************** Bottom of data **************************** |
```

Figure 11-4 Operating System Configuration: Work with Esoterics

8. In the command line, enter `add` (see Figure 11-5) to add an Esoteric.

```
+----------------------------- Esoteric List -----------------------------+
|   Goto  Filter  Backup  Query  Help                                     |
| - +--------------------- Add Esoteric ---------------------+ --------- |
|   |                                                         |           |
| C |                                                         | ==> CSR   |
|   | Specify the following values.                           |           |
| S |                                                         | 11.       |
|   | Esoteric name  . . . ITSOESO1                           |           |
| C | VIO eligible . . . . Yes    (Yes or No)                 |           |
| E | Token  . . . . . . . 1___                               |           |
|   |                                                         |           |
| / |                                                         |           |
| * |  F1=Help    F2=Split   F3=Exit    F5=Reset   F9=Swap    | ********** |
|   | F12=Cancel                                              |           |
|   +---------------------------------------------------------+           |
|                                                                         |
```

Figure 11-5 Operating System Configuration: Add Esoteric

9. Complete the following updates and press Enter:
   - Update Esoteric name to ITSOESO1.
   - Update `VIO eligible` to `Yes` (if you want to make this Esoteric VIO eligible).
   - Update `Token` to `1`.
10. After you add devices to the OSCONFIG ITSOTEST in the following examples, you then can add those device definitions to an Esoteric or Console, if applicable.

11.2.2 Defining more operating system LPARs in a CSS

A new (unreserved) partition includes the following considerations:

- 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. First, it must be redefined as reserved (Partition name = *). Next, the IODF must be activated on the processor, redefined to the new name, and then, the IODF is activated on the processor.

To change a reserved partition to an active partition in a 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 1 for this example), and press Enter.

5. Complete the following updates (see Figure 11-6) and press Enter:
   - Update `Partition Name` to MUSCA21 (a naming standard that is based on CSS=2, Partition =1).
   - Review `Partition usage` and change, if required. We use `OS` in this example.
   - Update `Description` to MUSCA21 test OS partition.

```
+----------------------------- Partition List -----------------------------+
|   Goto  Backup  Query  Help                                               |
| ------------------------------------------------------------------------ |
|                                                                          |
| Command ===> ________________________________________ Scroll ===> CSR    |
|                                                                          |
| Select one or more partitions, then press Enter. To add, use F11.        |
|                                                                          |
| Processor ID  . . . . : CETUS     CETUS                                  |
| Configuration mode  . : LPAR                                             |
| Channel S +-------------- Change Partition Definition ---------------+   |
|           |                                                          |   |
| / Partiti |                                                          |   |
| c *       | Specify or revise the following values.                  |   |
| _ *       |                                                          |   |
| _ *       | Partition name . . . MUSCA21                             |   |
| _ *       | Partition number . . 1      (same as MIF image ID)       |   |
| _ *       | Partition usage  . . OS     +                            |   |
| _ *       | UID uniqueness . . . N        (Y/N)                      |   |
| _ *       |                                                          |   |
|  F1=Help  | Description  . . . . MUSCA21 test OS partition______     |   |
|  F7=Backw |                                                          |   |
| F12=Cance |                                                          |   |
+---------- |  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset    | s-+
            |  F9=Swap   F12=Cancel                                    |
            +----------------------------------------------------------+
```

Figure 11-6 Processors: Change Partition Definition

12

Adding storage devices

This chapter describes how to define FICON switches and FICON Channel path IDs (CHPIDs), and connect them to a DASD Control Unit. It includes a list of these potential configuration items and a short description of how to complete each item by using the Hardware Configuration Definition (HCD) tool.

Note: Not all of the configuration items that are described in this chapter are necessarily required for your installation. Also, the examples that are presented are not intended to be exhaustive.

This chapter shows only the definition process. For more information about how to use HCD and FICON, see the following publications:

- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804
- *FICON Planning and Implementation Guide*, SG24-6497

This chapter includes the following topics:

12.1 Defining more I/O using HCD

When defining new I/O components in an IODF, certain definitions, such as OSCONFIGs, partitions, FICON Switches, Control Units, and Devices, must be made up front. After these items are defined, the following connections can be made:

- Defining FICON switches (Directors, SANs, and SAN switches)
- Defining FICON CHPIDs and connecting them to a FICON switch
- Defining FICON CHPIDs for direct connect to a 2107 Control Unit
- Defining FICON CHPIDs for switched connect to a 2107 Control Unit
- Defining 3390B devices to an OSCONFIG and EDT/Esoteric

The I/O definitions that are described next use HCD to demonstrate the examples. The examples continue by using the example that was created in Chapter 5, “Production IODF and setting up the CPC” on page 85 (for example, `ITSO.IODF78.WORK`).

12.1.1 Defining FICON switches (Directors, SANs, and SAN switches)

A new FICON switch includes the following considerations:

- 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 command line, enter `add` (see Figure 12-1 on page 277) to add a switch.
3. Complete the following updates and press Enter:
   - `Switch ID` to 41
   - `Switch type` to 2032
   - `Description` to ITSO test SAN definition
   - `Switch address` to 41
   - `Installed port range` to 00 and FE
   - `Switch CU number` to 0041
   - `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  (00-FF) 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  ____  ____  ____  ____              |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap     |
| F12=Cancel                                                          |
+---------------------------------------------------------------------+
```

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 control units

This section describes defining FICON CHPIDs, and then connecting them to FICON switches and a DASD control unit.

12.2.1 Defining FICON CHPIDs and connecting them to a FICON switch

A new FICON CHPID includes the following considerations:

- Channel path ID (CHPID).
- Channel ID (CHID).
- Channel path type.
- Operational mode.
- Description.
- Dynamic entry switch ID.
- Entry switch ID.
- Entry port.
- Partition access list.
- For performance and redundancy, how many CHPIDs are required to connect to the FICON Switch, and then to one or more Control Units?
- For performance and redundancy, how many I/O cards of that feature are installed in the processor and to what PCIe ports and to which CPC drawer do the I/O cards connect (see the PCHID/CHID report for a list of installed hardware)?
- After a CHPID is defined, it can then be:
  - Added to a predefined partition in that CSS
  - Connected to a FICON switch
  - Connected to a Control Unit

To define a new FICON CHPID, connect to a FICON switch, 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 12-3 on page 279) to add a CHPID.
5. Complete the following updates and press Enter:
   - `Channel path ID` to `14`
   - `Channel ID` to `114`
   - `Channel path type` to `FC`
   - `Operational mode` to SHR
   - `Description` to whatever description that you want
   - `Dynamic entry switch ID` to `41`
   - `Entry switch ID` to `41` (optional, but preferred)
   - `Entry port` to `10` (optional, but preferred)

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 2                                                  |
|                                                                           |
| 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#0428 16Gb FICON Exp16S+ 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  +                                            |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap      |
| F12=Cancel                                                                |
+---------------------------------------------------------------------------+
```

Figure 12-3 Processors: Add Channel Path

HCD now prompts you to select to which partition the CHPID can access.

6. Enter a forward slash '/' next to the partition that 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 : 2                                                  |
| Channel path ID  . . : 14     Channel path type  . : FC                   |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 2      MUSCA21          1      OS    MUSCA21 test OS partition          |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|  F1=Help       F2=Split      F3=Exit       F5=Reset      F6=Previous      |
|  F7=Backward   F8=Forward    F9=Swap      F12=Cancel                      |
+---------------------------------------------------------------------------+
```

Figure 12-4 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 and shows you 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 . . . . : MUSCA       Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 14    114   FC    SHR   41  41 10       No  FC#0428 16Gb FICON Exp16S+ SX
******************************* Bottom of data *****************************
```

Figure 12-5 Processors: Channel Path List

12.2.2 Defining FICON CHPIDs for direct connect to a 2107 Control Unit

A new FICON CHPID includes the following considerations:

- Channel path ID (CHPID).
- Channel ID (CHID).
- Channel path type.
- Operational mode.
- Description.
- Partition access list.
- For performance and redundancy, how many CHPIDs are required to connect to the FICON Switch and then to one or more Control Units?
- For performance and redundancy, 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 (see the PCHID/CHID report for a list of installed hardware)?
- After a CHPID is defined, it can then be added to a predefined partition in that CSS.
- The CHPID can then be connected to a Control Unit.

Note: For FC#0428 FICON Express16S+, defining FC and FCP CHPID types on the same I/O card is not supported. HCD issues the following warning message during a Validate or Build Production for an IODF:

```
W   CBDA963I Chpid type mix detected on processor MUSCA for channels:
             0.29, 0.31, 0.30, 0.44
```

To define a new FICON CHPID that connects directly to a Control Unit and provides 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.

5. Complete the following updates (see Figure 12-6) and press Enter:
   - `Channel path ID` to `15`
   - `Channel ID` to `115`
   - `Channel path type` to `FC`
   - `Operational mode` to `SHR`
   - `Description` to the description that you want

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 2                                                  |
|                                                                           |
| 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#0428 16Gb FICON Exp16S+ SX___                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap      |
| F12=Cancel                                                                |
+---------------------------------------------------------------------------+
```

Figure 12-6 Processors: Add Channel Path

6. Enter a forward slash '/' next to the partition that you want CHPID to have access to (see Figure 12-7), 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 : 2                                                  |
| Channel path ID  . . : 15     Channel path type  . : FC                   |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 2      MUSCA21          1      OS    MUSCA21 test OS partition          |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
|  F1=Help       F2=Split      F3=Exit       F5=Reset      F6=Previous      |
|  F7=Backward   F8=Forward    F9=Swap      F12=Cancel                      |
+---------------------------------------------------------------------------+
```

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 that shows 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 . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 14    114   FC    SHR   41  41 10       No  FC#0428 16Gb FICON Exp16S+ SX
_ 15    115   FC    SHR   __  __ __       No  FC#0428 16Gb FICON Exp16S+ SX
******************************* Bottom of data *****************************
```

Figure 12-8 Processors: Channel Path List

12.2.3 Defining FICON CHPIDs for switched connect to a 2107 Control Unit

FICON connections can be defined to a Control Unit by using one of two methods: through a FICON switch or by using direct connect.

The direct connect method mainly is used in an environment where only one processor is present as opposed to FICON switch connection in which multiple processors must connect to the same control units. This situation might not always be the case, however.

Connecting a FICON CHPID to a DASD Control Unit and its devices includes the following considerations:

- For performance and redundancy, how many CHPIDs are required to connect to the FICON Switch and then to one or more Control Units?
- For FICON Switched connections, a minimum of two FICON switches often are used that the FICON CHPIDs connect through, primarily for failure or service redundancy of the FICON Switches.

In our example, we connect to a predefined 2107 Control Unit (A000), with a CUADD of 40 and devices A000-A0EF (3390B) and A0F0-A0FF (3390A).

Note: For FC#0428 FICON Express16S+, defining FC and FCP CHPID types on the same I/O card are not supported. HCD issues the following warning message during a Validate or Build Production for an IODF:

```
W   CBDA963I Chpid type mix detected on processor MUSCA for channels:
             0.29, 0.31, 0.30, 0.44
```

To define FICON CHPIDs for switched connect, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit you want to connect to or in the command Line, enter `L A000`. In our example, we use `A000`.

3. Enter `c` next to the Control Unit definition, and press Enter.
4. Make the following updates for a FICON switch connection (see Figure 12-9), and press Enter:
   – Update `Connected to switches` to `41`. Switch `41` is the switch ID that we defined in the previous example.
   – Update `Ports` to `11`. Port `11` is the FICON switch port that the Control Unit 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  . . . . . . . . DS8886 - SU2 - SSID=A000             |
|                                                                    |
| Connected to switches  . . . 41  __ __ __ __ __ __ __  +          |
| Ports  . . . . . . . . . . . 11  __ __ __ __ __ __ __  +          |
|                                                                    |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                    |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap   |
| F12=Cancel                                                        |
+--------------------------------------------------------------------+
```

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 Control Unit 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 control unit and also has access to the CHPID that we want to connect to the control unit. In our example, we use MUSCA.2. Then, press Enter.

6. Complete the following updates to define the Processor / CU connection panel (see Figure 12-10), and press Enter:
   - `Channel path IDs` to `14`. Channel path ID 14 is the CHPID that we defined in the previous example.
   - `Link address` to `4111`. Link address 4111 is FICON switch 41 and Control Unit Port 11.
   - `Unit address` to `00`. Unit address of 00 is the starting unit address number on the CHPID.
   - `Number of units` to `256`. Number of units of 256 equals A000-A0FF (that is, 00-FF or 256).
   - `Logical address` to `40`. Logical address of 40 is the CUADD or the Control Unit Address that is defined in the DS8886 that defines the location of the devices in the DS8886.

```
                             Select Processor / CU
+---------------------- Change Control Unit Definition -----------------------+
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Control unit number  . : A000        Type . . . . . . : 2107                 |
| Processor ID . . . . . : MUSCA       Musca                                   |
| Channel Subsystem ID . : 2                                                   |
|                                                                              |
| 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)                                   |
|                                                                              |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap         |
| F12=Cancel                                                                   |
+------------------------------------------------------------------------------+
```

Figure 12-10 Control units: Change Control Unit Definition detail

HCD now shows the Modify Device Parameters panel in which you can override the Unit Address (UA) numbers. For a 2107 DASD Control Unit definition, the starting UA often is `00` and the ending UA is `FF`, which gives you 256 DASD device definitions for the Control Unit.

7. In our example, we do not change the defaults that were proposed by HCD (see Figure 12-11). Press Enter.

```
                         Select Processor / CU
+------------------------- Modify Device Parameters -------------------------+
|                                                   Row 1 of 2 More:       > |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Specify or revise any changes to the device parameters in the list below.  |
| To view attached control units, scroll to the right.                       |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                  |
| Channel Subsystem ID : 2                                                   |
|                                                                            |
| ---------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         0   F0  F0    No       Yes    __                  |
| ***************************** Bottom of data ***************************** |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt      F5=Reset         |
|  F7=Backward   F8=Forward    F9=Swap      F12=Cancel     F20=Right         |
| F22=Command                                                                |
+----------------------------------------------------------------------------+
```

Figure 12-11   Control units: Modify Device Parameters

HCD now returns to the Select Processor / CU panel, which shows the CHPID (14) to Destination Link Address (11) through the FICON switch (41) connection definition (see Figure 12-12).

```
                            Select Processor / CU     Row 1 of 9 More:
Command ===> _______________________________________________ Scroll ===> CSR

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------
_ MUSCA.2    14.4111 _______ _______ _______ _______ _______ _______ _______
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    _______ _______ _______ _______ _______ _______ _______ _______
******************************* Bottom of data ****************************
```

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 Control Unit definition, you omit the Link address values.

8. Update `Channel path IDs` to `15` to define the Processor/CU connection (see Figure 12-13 on page 287), and press Enter.

```
                              Select Processor / CU
+---------------------- Change Control Unit Definition -----------------------+
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                     |
|                                                                              |
| Control unit number  . : A000        Type . . . . . . : 2107                |
| Processor ID . . . . . : MUSCA       Musca                                  |
| Channel Subsystem ID . : 2                                                  |
|                                                                              |
| 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

9. HCD again shows the Modify Device Parameters panel in which you can override the Unit Address (UA) numbers. Press Enter to continue.

   HCD now returns to the Select Processor / CU panel that shows the CHPID (15) connection definition (see Figure 12-14).

```
                            Select Processor / CU     Row 1 of 9 More:
Command ===> _______________________________________________ Scroll ===> CSR

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------
_ MUSCA.2    14.4111 15      _______ _______ _______ _______ _______ _______
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    _______ _______ _______ _______ _______ _______ _______ _______
******************************* Bottom of data ******************************
```

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 Control Unit, this configuration is possible.

A typical scenario might be if you were moving from direct connected DASD to FICON switch connected DASD, but you cannot take the DASD offline to live systems.

12.2.4 Defining 3390B devices to an OSCONFIG and EDT/Esoteric

The OSCONFIG name is the part of an IODF that determines which devices a z/OS system can access to when it IPLs. Also, the partition that the z/OS system is restarted in must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups and 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 the UNIT = parameter is used in a **`JCL DD`** statement. However, this allocation can be overridden or intercepted by using DFSMS.

The OSCONFIG name includes the following items:

- EDT ID (Eligible Device Table ID): Esoterics / VIO
- Consoles

Adding devices to an OSCONFIG and Esoteric includes the following considerations:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access that device.
- Does the device need to be added to an Esoteric device group? Mostly, this issue 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 you want to add to the OSCONFIG or, in the command line 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 Control Unit that the devices are attached to (see Figure 12-15 on page 289). 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.                                                          |
|                                                                             |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap        |
| F12=Cancel                                                                  |
+-----------------------------------------------------------------------------+
```

Figure 12-15 I/O Devices: Change Device Definition

HCD now displays the Device / Processor Definition panel in which you can modify some of the Device parameters that relate to SubChannel Set (SS), Unit Address (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 ===> CSR   |
|                                                                           |
| 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 MUSCA.2     _    00   No        Yes     __         No             ___   |
| ***************************** Bottom of data **************************** |
|                                                                           |
|                                                                           |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset         |
|  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel        |
| F22=Command                                                               |
+---------------------------------------------------------------------------+
```

Figure 12-16 I/O Devices: Device / Processor Definition continued

5. Next is the HCD panel in which you define devices to the operating system configuration. 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 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 specific device type. In our example, we are adding 3390B and 3390A devices to ITSOTEST.
7. Complete the following updates to define the Processor/CU connection (see Figure 12-17), and press Enter:
   - `OFFLINE` to `No` (if you want these devices to be Online during IPL time)
   - `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically)
   - `LOCANY` to `Yes` (if the device UCB can be in 31-bit storage)
   - `WLMPAV` to `Yes` (if you want the device to be managed by Work Load Manager)
   - `READ-ONLY` to `No` (use to set Secondary devices to read only)
   - `SHARED` to `Yes` (if the device is going to be shared between other systems)
   - `SHAREDUP` to `No` (needs to be set to `No` is Shared is set to `Yes`)

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                 Row 1 of 7 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise the values below.                                         |
|                                                                             |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1            |
| 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 ***************************** |
|                                                                             |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt      F5=Reset        |
|  F7=Backward   F8=Forward    F9=Swap      F12=Cancel     F22=Command      |
+-----------------------------------------------------------------------------+
```

Figure 12-17 I/O Devices: Define Device Parameters / Features

The Assign/Unassign Device to Esoteric panel is now displayed in which you can specify to which Esoteric (if any) you want the devices to be added.

8. Under `Assigned`, enter `Yes` to update the Assign/Unassign the Device to Esoteric (see Figure 12-18), and press Enter.

```
+-------------------- Assign/Unassign Device to Esoteric --------------------+
|                                                                Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR   |
|                                                                           |
| 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 z14 ZR1       |
| 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 ***************************** |
|                                                                           |
|                                                                           |
|  F1=Help        F2=Split       F3=Exit        F4=Prompt      F5=Reset     |
|  F6=Previous    F7=Backward    F8=Forward     F9=Swap       F12=Cancel    |
| F22=Command                                                               |
+---------------------------------------------------------------------------+
```

Figure 12-18 I/O Devices: Assign/Unassign Device to Esoteric

9. The final panel is now displayed that 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 34 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : A000          Number of devices  : 240             |
| Device type  . . : 3390B                                              |
|                                                                       |
| / Config. ID   Type     SS Description                     Defined    |
| _ DBSV4SU4     MVS         z/OS DB Server 4                           |
| _ DBSV5SU4     MVS         z/OS DB Server 5                           |
| _ DBSV6SU4     MVS         z/OS DB Server 6                           |
| _ ITSO         MVS         All ITSO devices                           |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes       |
| _ PERF4SU4     MVS         z/OS Appl Server 4                         |
| _ PERF5SU4     MVS         z/OS Perf Server 5                         |
| _ PERF6SU4     MVS         z/OS Appl Server 6                         |
| _ ZHLXDRVR     MVS         z/OS Helix Central Mgmt Image              |
| _ ZMEDRVR      MVS         z/OS 2.1 ServerPak refresh                |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset     |
|  F6=Previous   F7=Backward   F8=Forward    F9=Swap      F12=Cancel    |
| F22=Command                                                           |
+-----------------------------------------------------------------------+
```

Figure 12-19 I/O Devices: Assign/Unassign Device to Esoteric completed

The same steps can be used for the remaining devices A0F0-A0FF (3390A) for this example.

13

Adding network devices

This chapter describes how to define OSC, OSD, OSE, OSM, and IQD CHPIDs; Control Units; and devices.

It includes a list of these potential configuration items and a short description of how to complete each item by using the Hardware Configuration Definition (HCD) tool.

> **Note:** Not all of the configuration items that are described in this chapter are necessarily required for your installation. Also, the examples that are presented are not intended to be exhaustive.

This chapter includes the following topics:

13.1 Defining more I/O using HCD

When defining new I/O components in an IODF, specific definitions, such as OSCONFIGs, partitions, FICON Switches, Control Units, and Devices, must be made up front. After the following items are defined, the connections can be made:

- OSC CHPIDs
- OSC CHPID connections to an OSC Control Unit
- 3270-X devices to an OSCONFIG
- 3270-X devices to NIP within an OSCONFIG
- OSD CHPIDs
- OSD CHPID connections to an OSA Control Unit
- OSA and OSAD devices to an OSCONFIG
- OSE CHPIDs
- OSE CHPID connections to an OSA Control Unit
- OSA and OSAD devices to an OSCONFIG
- OSM CHPIDs
- OSM CHPID connections to an OSM Control Unit
- OSA-M devices to an OSCONFIG
- IQD CHPIDs
- IQD CHPID connections to an IQD Control Unit
- IQD devices to an OSCONFIG

The I/O definitions that are described next use the HCD tool to demonstrate the examples. The examples continue by using the example that was created in Chapter 5, “Production IODF and setting up the CPC” on page 85 (for example, `ITSO.IODF78.WORK`).

13.2 OSA CHPID definitions

This section describes defining OSC, OSD, OSE, and OSM CHPIDs and their control units and devices.

13.2.1 Defining OSC CHPIDs

When defining an OSD or OSE connection, you first must determine which type of OSA card you need for your configuration. OSC connections use the copper-based OSA Express6S 1000BASE-T card:

- Optical connection cards:
  - OSA Express6S GbE SX or LX
  - OSA Express6S 10 Gb SR or LR
- Copper-based connection cards: OSA Express6S 1000BASE-T

A new OSC CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Partition access list

- For performance and redundancy, how many I/O cards of that feature are installed in the processor and to what PCIe ports on what CPC drawer that the I/O cards connect to (see the PCHID/CHID report for a list of installed hardware)

To define a new 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 CSS ID that you want to add a CHPID to, and press Enter.
4. In the command line, enter `add` (see Figure 13-1) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 1                                                  |
|                                                                           |
| Channel path ID  . . . . BO    +            Channel ID  110  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSC   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . FC#0426 OSA Express6S 1000Base-T                 |
|                                                                           |
| 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

5. Complete the following updates and press Enter:
   - `Channel path ID` to `BO`
   - `Channel ID` to `110`
   - `Channel path type` to `OSC`
   - `Operational mode` to `SHR`
   - `Description` to the description that you want

6. HCD prompts you to select to which partition the CHPID should have access. Enter a forward slash '/' next to the partition that you want (see Figure 13-2), 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 : 1                                                  |
| Channel path ID  . . : B0     Channel path type  . : OSC                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA1A          A      OS                                       |
| / 1      MUSCA1B          B      OS                                       |
| / 1      MUSCA1C          C      OS                                       |
| / 1      MUSCA1D          D      OS                                       |
| / 1      MUSCA1E          E      OS                                       |
| / 1      MUSCA1F          F      OS                                       |
| / 1      MUSCA11          1      OS                                       |
| / 1      MUSCA12          2      OS                                       |
| / 1      MUSCA13          3      OS                                       |
| / 1      MUSCA14          4      OS                                       |
| / 1      MUSCA15          5      OS                                       |
| / 1      MUSCA16          6      OS                                       |
*---------------------------------------------------------------------------*
```

Figure 13-2 Processors: Define Access List, OSC

Because multiple partitions were defined in this CSS, we see the Define Candidate List panel, although we defined the CHPID as SHR. We do not connect any partitions to this CHPID from the candidate list. Instead, press `Enter` on this panel.

HCD now returns back to the Channel Path List, which shows you the CHPID that was just defined (see Figure 13-3).

```
                              Channel Path List     Row 10 of 35 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 1

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 48    15C   FC    SPAN  01  01 1A       No  Disk *** was 01.24
_ 49    15D   FC    SPAN  02  02 1C       No  Disk *** was 02.24
_ 54    17D   FC    SPAN  01  01 0C       No  FCTC
_ 55    179   FC    SPAN  02  02 0C       No  FCTC
_ 58    13C   FC    SPAN  01  01 1B       No  Disk *** was 01.25
_ 59    13D   FC    SPAN  02  02 1D       No  Disk *** was 02.25
_ B0    110   OSC   SHR   __  __ __       No  FC 0426 OSA Express6S 1000Base-T
```

Figure 13-3 Processors: Channel Path List, OSC

13.2.2 Defining OSC CHPID connections to an OSC Control Unit

The only way to define an OSA connection to its Control Unit is direct connected.

You might want to connect the OSC Control Unit definition to multiple CPCs even though the physical OSC is still unique to any one CPC. Also, you might want to span the OSC over multiple CSSs within a CPC.

Consider the following points for connecting an OSC CHPID to an OSC Control Unit and its 3270-X devices:

- How many OSCs are required to provide a primary and secondary/backup network connection.
- In our example, we connect to a predefined OSC Control Unit (1B10) and 3270-X devices 1B10-1B1F.

To define OSC CHPID connections to an OSC Control Unit, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit that you want to connect to or in the command line enter `L 1B00`. In our example, we use `1B00`.
3. Enter `c` next to the Control Unit definition, and press Enter.
4. Make any changes to the Control Unit definition that you want (see Figure 13-4), and press Enter.

```
+---------------- Change Control Unit Definition ------------+
|                                                            |
|                                                            |
| Specify or revise the following values.                    |
|                                                            |
| Control unit number  . . . . 1B00  +                       |
|                                                            |
| Control unit type  . . . . . OSC           +               |
|                                                            |
| Serial number  . . . . . . . __________  +                 |
| Description  . . . . . . . . ________________________________ |
|                                                            |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +    |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +    |
|                                                            |
| Define more than eight ports . . 2   1.  Yes               |
|                                      2.  No                |
|                                                            |
|                                                            |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap |
| F12=Cancel                                                 |
+------------------------------------------------------------+
```

Figure 13-4 Control units: Change Control Unit Definition, OSC

HCD now shows the Select Processor / CU panel. This panel is where the connection is now made between the OSC CHPID (B0) and the Control Unit (1B00).

5. Enter `c` next to the Processor.CSS that contains the partition we want to have access to the control unit and also has access to the CHPID we want to connect to the control unit. Then, press Enter.

6. Update `Channel path IDs` to `B0` (channel path ID B0 is the CHPID we defined in the previous example) to define the Processor / CU connection (see Figure 13-5), and press Enter.

```
*---------------------- Change Control Unit Definition -----------------------*
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Control unit number  . : 1B00        Type . . . . . . : OSC                  |
| Processor ID . . . . . : MUSCA       Musca                                   |
| Channel Subsystem ID . : 1                                                   |
|                                                                              |
| Channel path IDs . . . . B0    __    __    __    __    __    __    __    +   |
| 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

HCD now shows the Modify Device Parameters panel where you can override the Unit Address (UA) numbers. For most OSA definitions, the Unit address starts at `00`.

7. Update `UA New` to `00` to define the Modify Device Parameters (see Figure 13-6), and press Enter.

```
*------------------------- Modify Device Parameters -------------------------*
|                                                        Row 1 of 1 More:      > |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise any changes to the device parameters in the list below.  |
| To view attached control units, scroll to the right.                       |
|                                                                             |
| Processor ID . . . . : MUSCA        Musca                                   |
| Channel Subsystem ID : 1                                                    |
|                                                                             |
| ---------Device---------     ---UA----                Preferred Exposure    |
| No., Range Type          SS+ Old New + Time-Out STADET CHPID +   Device     |
| 0800,016   3270-X        _   00  00    No       No     __                   |
| ***************************** Bottom of data ***************************** |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 13-6 Control units: Modify Device Parameters, OSC

HCD now returns back to the Select Processor / CU panel that shows the CHPID (B0) connection definition (see Figure 13-7).

```
 Select Processor / CU    Row 1 of 12 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : 1B00     Control unit type . . . : OSC

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    B0      _______ _______ _______ _______ _______ _______ _______
_ MUSCA.2    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-7 Processors: Select Processor / CU: CHPID to Link address connection, OSC

8. By pressing **F20 (Right)**, you can see the other parts of the definition summary.

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 IPLs. The partition that the z/OS system is restarted in also must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups, which are defined in EDTs within an OSCONFIG. OSA definitions often do not use Esoterics.

Consider the following points 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 our example, we add the predefined OSC devices 0800-080F (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 0800`. In our example, we use `1B10`.
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 Control Unit that the devices are attached to (see Figure 13-8). Press Enter.

```
+------------------------ Change Device Definition -------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Device number  . . . . . . . : 0800 (0000 - FFFF)                         |
| Number of devices  . . . . . : 16                                         |
| Device type  . . . . . . . . : 3270-X                                     |
|                                                                           |
| Serial number  . . . . . . . . __________  +                              |
| Description  . . . . . . . . . ________________________________           |
|                                                                           |
| Volume serial number . . . . . ______  + (for DASD)                       |
|                                                                           |
| PPRC usage . . . . . . . . . . _  + (for DASD)                            |
|                                                                           |
| Connected to CUs . 1B00  ____  ____  ____  ____  ____  ____  ____  +      |
|                                                                           |
| ENTER to continue.                                                        |
|                                                                           |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap      |
| F12=Cancel                                                                |
+---------------------------------------------------------------------------+
```

Figure 13-8 I/O Devices: Change Device Definition, OSC

5. HCD now displays the Device / Processor Definition panel where you can modify some of the Device parameters that relate to SubChannel Set, Unit Address, and Explicit Device Candidate List (see Figure 13-9).

```
*---------------------- Device / Processor Definition -----------------------*
|                                                                 Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Select processors to change device/processor definitions, then press       |
| Enter.                                                                     |
|                                                                            |
| Device number  . . : 0800       Number of devices  . : 16                  |
| Device type  . . . : 3270-X                                                |
|                                                                            |
|                                          Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null  |
| _ MUSCA.1     _    00   No        No      __         No             ___   |
| ***************************** Bottom of data ***************************** |
|                                                                            |
```

Figure 13-9 I/O Devices: Device / Processor Definition continued, OSC

6. Next is the HCD panel where we Define Devices to the Operating System Configuration. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG you want to add to the Devices to, or in the command line enter `L ITSOTEST`. In our example, we use ITSOTEST.
7. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the Device Parameters and Features applicable to that particular device type. In our example, we are adding 3270-X devices to ITSOTEST.
8. Complete the following updates to define the Processor / CU connection (see Figure 13-10 on page 302), and press Enter:
   - `OFFLINE` to `No` (if you want these devices to be Online during IPL time).
   - `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   - `LOCANY` to `Yes` (if the device UCB can be in 31-bit storage).
   - `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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1             |
| Device number  . . : 0800         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 13-10 I/O Devices: Define Device Parameters / Features, OSC

9. The Assign/Unassign Device to Esoteric panel is now displayed in which you can specify to which Esoteric (if any) you want the devices to be added. Press Enter (see Figure 13-11 on page 303).

   We add only the OSC/3270-X devices to the OSCONFIG ITSOTEST and not to any esoterics in our example.

```
*----------- Define Device to Operating System Configuration -----------*
|                                                              Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                        |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                        |
| Device number  . : 0800          Number of devices  : 16              |
| Device type  . . : 3270-X                                              |
|                                                                        |
| / Config. ID   Type     SS Description                     Defined    |
| _ COD230       MVS         z/OS image for COD test         Yes        |
| _ DBSV4SU4     MVS         z/OS DB Server 4                Yes        |
| _ DBSV5SU4     MVS         z/OS DB Server 5                Yes        |
| _ DBSV6SU4     MVS         z/OS DB Server 6                Yes        |
| _ ITSO         MVS         All ITSO devices                Yes        |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes        |
| _ PERF4SU4     MVS         z/OS Appl Server 4              Yes        |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                Yes        |
| _ PERF6SU4     MVS         z/OS Appl Server 6              Yes        |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh     Yes        |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system           Yes        |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system          Yes        |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image            Yes        |
*-----------------------------------------------------------------------*
```

Figure 13-11 I/O Devices: Assign/Unassign Device to Esoteric, OSC

10. The final panel is now displayed in which the devices that are defined to the OSCONFIG are shown. Press Enter to return to the I/O Device List.

13.2.4 Defining 3270-X devices to NIP within an OSCONFIG

The NIP Console list determines the device addresses that are eligible to receive nucleus initialization program (NIP) or IPL messages in the early startup stages of when z/OS is started.

The devices first must 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, NIP tries the next device in the list until all devices in the list are tried.

If NIP cannot write IPL messages to any 3270-X device in the list, the messages are written to the HMC Operating System Messages panel.

To view these messages, complete the following steps:

1. Select the LPAR for IPL on the HMC.
2. Click the >> breakout symbol that is next to the LPAR name.

3. Select **Daily** → **Operating System Messages**, as shown in Figure 13-12.

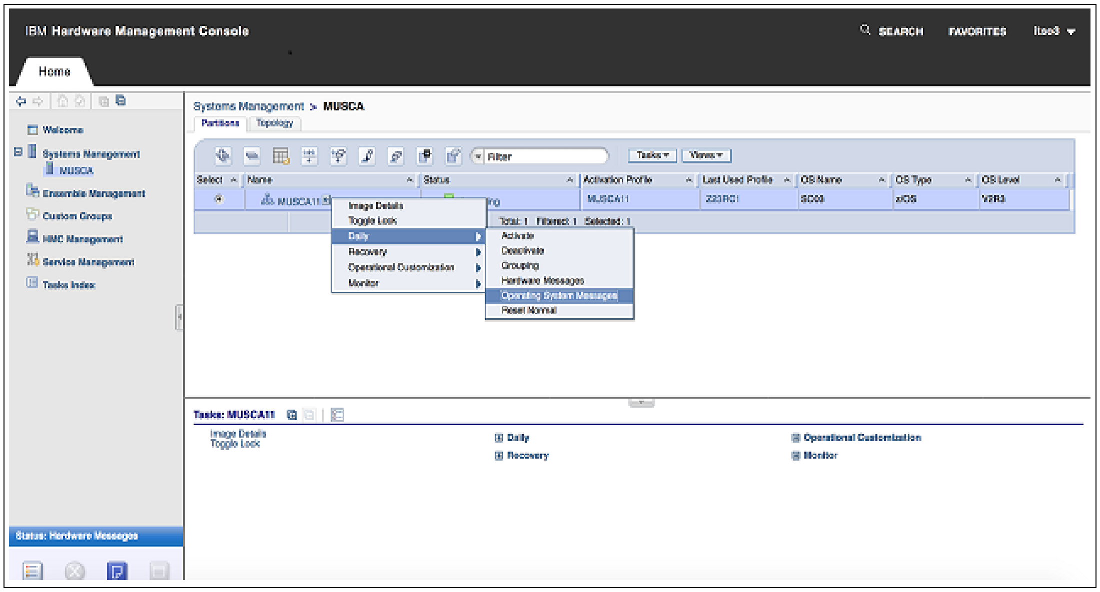

Figure 13-12 Operating System Messages

Commands and displays can be entered in the Command box (see Figure 13-13).

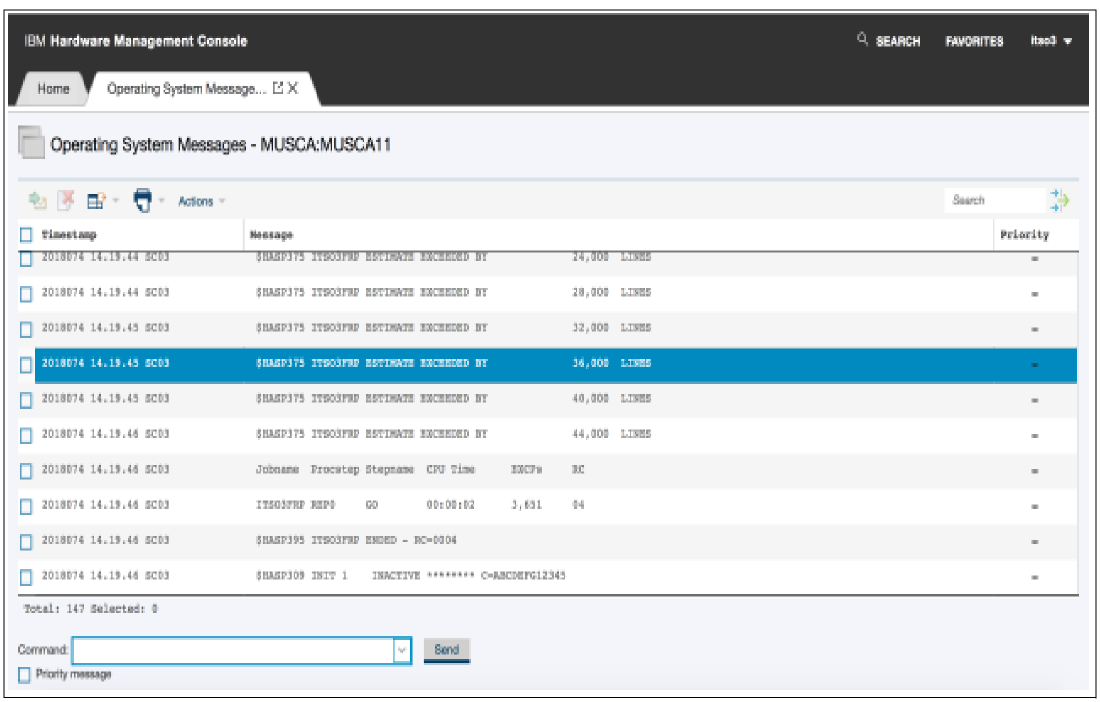

Figure 13-13 Operating System Messages command interface

Consider the following points for adding devices to a NIP console list within an OSCONFIG:

- Adding a device to a NIP console list within an OSCONFIG does not necessarily mean that NIP can write IPL messages to that device.
- The devices that are defined in the NIP console list also need Control unit and CHPID access to the partition into where z/OS is being started.

- On the HMC, under OSA Advanced Facilities, the OSC (OSA-ICC) console Server and Session definitions also must be defined and activated.
- A valid 3270-X session (that uses IBM PCOM or an equivalent 3270 emulator) also must be connected to the OSA-ICC Session. This configuration allows a valid session to be established to the OSA-ICC for NIP messages to be delivered to that device.
- Our example adds the predefined OSC devices 1B10-1B11 (3270-X).

To define 3270-X devices to 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 ITSOTEST`. In our example, we use ITSOTEST.
3. Enter `n` next to the OSCONFIG, and press Enter.

   HCD displays the currently defined devices in the NIP Console List (see Figure 13-14).

```
*--------------------------- NIP Console List ----------------------------*
|   Goto  Backup  Query  Help                                              |
| ----------------------------------------------------------------------  |
|                                                                          |
| Command ===> _______________________________________ Scroll ===> CSR     |
|                                                                          |
| Select one or more consoles, then press Enter. To add, use F11.          |
|                                                                          |
| Configuration ID . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1         |
|                                                                          |
|   Order   Device                                                         |
| / Number  Number  Device Type                                            |
| *************************** Bottom of data **************************** |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
|                                                                          |
*--------------------------------------------------------------------------*
```

Figure 13-14 OCSONFIGs: NIP Console List

In our example, no devices are defined in the NIP Console List.

4. In the command line, enter `add` (see Figure 13-15) to add a 3270-X device to the NIP Console List.
5. Update `Device number of console` to `0800` 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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1        |
|                  *-------------- Add NIP Console ---------------*       |
|   Order   Device |                                              |       |
| / Number  Number |                                              |       |
| **************** | Specify the following values.                | ***** |
|                  |                                              |       |
|                  | Device number of console . . . . . . 0800    |       |
|                  |                                              |       |
|                  | Order number . . . . . . . . . . . . 1       |       |
|                  |                                              |       |
|                  |                                              |       |
|                  |                                              |       |
|                  *----------------------------------------------*       |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 13-15 OSCONFIGs: Add NIP Console

Because this is the first device entry in the list, the order is 1 (see Figure 13-16).

```
*--------------------------- 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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1        |
|                                                                         |
|   Order   Device                                                        |
| / Number  Number  Device Type                                           |
| _   1     0800    3270-X                                                |
| *************************** Bottom of data **************************** |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 13-16 OCSONFIGs: NIP Console added

6. Add device 0801 to the NIP Console List (see Figure 13-17).

```
*--------------------------- 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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1        |
|                                                                         |
|   Order   Device                                                        |
| / Number  Number  Device Type                                           |
| _    1     0800    3270-X                                               |
| _    2     0801    3270-X                                               |
| *************************** Bottom of data **************************** |
|                                                                         |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 13-17 OCSONFIGs: extra NIP Console added

13.2.5 Defining OSD CHPIDs

When defining an OSD or OSE connection, first you must determine which type of OSA card you need for your configuration. OSD and OSE connections can use optical or copper-based cards. The type of card that is used depends on the type and speed of network to which you want to connect your OSAs. The following types of cards are available:

- Optical-based:
  - OSA Express6S GbE SX or LX
  - OSA Express6S 10 Gb SR or LR
- Copper-based: OSA Express6S 1000BASE-T

A new OSD CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Partition access list
- For performance and redundancy, 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 (see the PCHID/CHID report for a list of installed hardware)

To define a new 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 to which you want to add a CHPID. Press Enter.
3. Enter s next to the CSS ID to which you want to add a CHPID. Press Enter.

4. In the command line, enter `add` (see Figure 13-18) to add a CHPID.
5. Complete the following updates and press Enter:
   - `Channel path ID` to `B8`.
   - `Channel ID` to `148`.
   - `Channel path type` to `OSD`.
   - `Operational mode` to `SHR`.
   - `Description` to the description that you want.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 1                                                  |
|                                                                           |
| Channel path ID  . . . . B8    +            Channel ID  148  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSD   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0425 OSA Express6s 10GbE SR                   |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 13-18 Processors: Add Channel Path, OSD

6. HCD now prompts you to specify whether the allow more than 160 TCP/IP stacks. The default is `No`, which we use for our example.
7. Leave the default option for 160 TCP/IP stacks to `No` for our example (see Figure 13-19), and 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                     |
|  F1=Help     F2=Split    F3=Exit     F5=Reset    F9=Swap    |
| F12=Cancel                                                  |
+-------------------------------------------------------------+
```

Figure 13-19 Processors: Allow for more than 160 TCP/IP stacks, OSD

8. 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.
9. Leave the default option for Physical Network IDs to blank fields (see Figure 13-20), 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  . . ________________                              |
|                                                                          |
|                                                                          |
|                                                                          |
|  F1=Help     F2=Split    F3=Exit     F5=Reset    F9=Swap    F12=Cancel   |
+--------------------------------------------------------------------------+
```

Figure 13-20 Processors: Add/Modify Physical Network IDs, OSD

10. HCD prompts you to select to which partition the CHPID can access. Enter a forward slash '/' next to the partition that you want to use (see Figure 13-21), and press Enter.

```
*-------------------------- Define Access List ---------------------------*
|                                                             Row 4 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR  |
|                                                                          |
| Select one or more partitions for inclusion in the access list.         |
|                                                                          |
| Channel subsystem ID : 1                                                 |
| Channel path ID  . . : B8     Channel path type  . : OSD                 |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| / 1      MUSCA1D          D      OS                                      |
| / 1      MUSCA1E          E      OS                                      |
| / 1      MUSCA1F          F      OS                                      |
| / 1      MUSCA11          1      OS                                      |
```

Figure 13-21 Processors: Define Access List, OSD

Because we 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.

The HCD now returns to the Channel Path List panel, which shows you the CHPID that was just defined (see Figure 13-22).

```
 Channel Path List     Row 13 of 35 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 1

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 55    179   FC    SPAN  02  02 0C       No  FCTC
_ 58    13C   FC    SPAN  01  01 1B       No  Disk *** was 01.25
_ 59    13D   FC    SPAN  02  02 1D       No  Disk *** was 02.25
_ B0    110   OSC   SHR   __  __ __       No  FC 0426 OSA Express6S 1000Base-T
_ B1    130   OSD   SPAN  __  __ __       No  1000BaseT
_ B2    14C   OSD   SPAN  __  __ __       No  1000BaseT
_ B3    12C   OSD   SPAN  __  __ __       No  1000BaseT
_ B4    170   OSD   SPAN  __  __ __       No  1000BaseT
_ B5    150   OSD   SPAN  __  __ __       No  1000BaseT
_ B8    148   OSD   SHR   __  __ __       No  FC#0425 OSA Express6s 10GbE SR
```

Figure 13-22 Processors: Channel Path List, OSD

13.2.6 Defining OSD CHPID connections to an OSA Control Unit

The only way to define an OSA connection to its Control Unit is direct connected.

You might want to connect the OSA Control Unit 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.

Consider the following points when connecting an OSD CHPID to an OSA Control Unit and its OSA/OSAD devices:

- How many OSAs are required to provide a primary and secondary/backup network connection.
- Our example connects to a predefined OSA Control Unit (1B80) and OSA devices 1B80-1B8E and OSAD device 1B8F.

To define OSD CHPID connections to an OSA Control Unit, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit you want to connect to, or in the command line enter `L 1B80`. In our example, we use `1B80`.
3. Enter `c` next to the Control Unit definition, and press Enter.

4. Make any changes to the Control Unit definition that you want (see Figure 13-23), and press Enter.

```
*----------------- Change Control Unit Definition ------------------*
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . 1B80  +                              |
|                                                                   |
| Control unit type  . . . . . OSA           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . 10G Copper - MUSCA                   |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
*-------------------------------------------------------------------*
```

Figure 13-23 Control units: Change Control Unit Definition, OSD

HCD now shows the Select Processor / CU panel. It is in this panel that the connection is now made between the OSD CHPID (B0) and the Control Unit (1B80).

5. Enter c next to the Processor.CSS that contains the partition we want to have access to the control unit, and also has access to the CHPID we want to connect to the control unit. Then, press Enter.

6. Update `Channel path IDs` to `B8` (channel path ID B8 is the CHPID that we defined in the previous example) to define the Processor / CU connection (see Figure 13-24). Press Enter.

```
*---------------------- Change Control Unit Definition -----------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : 1B80        Type . . . . . . : OSA                 |
| Processor ID . . . . . : MUSCA       Musca                                  |
| Channel Subsystem ID . : 1                                                  |
|                                                                             |
| Channel path IDs . . . . B8    __    __    __    __    __    __    __    +  |
| 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 13-24 Control units: Change Control Unit Definition detail, OSD

7. HCD now shows the Modify Device Parameters panel where you can override the Unit Address (UA) numbers. For most OSA definitions, the Unit address starts at `00`.

8. Complete the following updates to define the Modify Device Parameters (see Figure 13-25), and press Enter:
   - `UA New` to `00` for the OSA device type
   - `UA New` to `FE` for the OSAD device type

```
*------------------------- Modify Device Parameters -------------------------*
|                                                   Row 1 of 2 More:       > |
| Command ===> __________________________________________ Scroll ===> CSR     |
|                                                                            |
| Specify or revise any changes to the device parameters in the list below.  |
| To view attached control units, scroll to the right.                       |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                  |
| Channel Subsystem ID : 1                                                   |
|                                                                            |
| ---------Device---------     ---UA----                 Preferred Exposure  |
| No., Range Type          SS+ Old New + Time-Out STADET CHPID +   Device     |
| 1B80,015   OSA           _   00  00    No       No     __                  |
| 1B8F,001   OSAD          _   FE  FE    No       No     __                  |
| **************************** Bottom of data ****************************** |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
```

Figure 13-25 Control units: Modify Device Parameters, OSD

HCD now returns to the Select Processor / CU panel, which shows the CHPID (B2) connection definition (see Figure 13-26).

```
Select Processor / CU    Row 1 of 12 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : 1B80     Control unit type . . . : OSA

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    B8      _______ _______ _______ _______ _______ _______ _______
_ MUSCA.2    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-26 Processors: Select Processor / CU: CHPID to Link address connection, OSD

9. Press F20 (Right) to see the other parts of the definition summary.

13.2.7 Defining OSA and OSAD devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines which devices a z/OS system can access when it IPLs. The partition that the z/OS system is started in also must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups, which are defined in EDTs within an OSCONFIG. OSA definitions often do not use Esoterics.

Adding devices to an OSCONFIG includes the following considerations:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access to that device.
- Our example adds the predefined OSA devices 1B80-1B8F (OSA/OSAD).

To define OSA and OSAD 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 1B80`. In our example, we use `1B80`.

   Because OSA and OSAD are two different device types, they must be added separately to the OSCONFIG.
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 Control Unit to which the devices are attached (see Figure 13-27). Press Enter.

```
*------------------------ Change Device Definition -------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number . . . . . . . : 1B80  (0000 - FFFF)                           |
| Number of devices . . . . . : 15                                            |
| Device type . . . . . . . . : OSA                                           |
|                                                                             |
| Serial number . . . . . . . . __________  +                                 |
| Description . . . . . . . . . ________________________________              |
|                                                                             |
| Volume serial number . . . . . ______  + (for DASD)                         |
|                                                                             |
| PPRC usage . . . . . . . . . . _  + (for DASD)                              |
|                                                                             |
| Connected to CUs . 1B80  ____  ____  ____  ____  ____  ____  ____  +        |
|                                                                             |
| ENTER to continue.                                                          |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 13-27 I/O Devices: Change Device Definition, OSD

5. HCD now displays the Device / Processor Definition panel where you can modify some of the Device parameters that relate to SubChannel Set, Unit Address, and Explicit Device Candidate List. Enter c next to the Proc.CCSID item if you want to modify any of these parameters (see Figure 13-28), 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  . . : 1B80       Number of devices  . : 15                  |
| Device type  . . . : OSA                                                   |
|                                                                            |
|                                          Preferred  Device Candidate List  |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null   |
| _ MUSCA.1     _    00   No        No      __         No             ___    |
| ***************************** Bottom of data ***************************** |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 13-28 I/O Devices: Device / Processor Definition continued, OSD

6. Next is the HCD panel where we Define Devices to the Operating System Configuration. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG you want to add to the Devices to, or in the command line enter `L ITSOTEST`. In our example, we use ITSOTEST.
7. Enter s next to the OSCONFIG, and press Enter.

   HCD displays the Device Parameters and Features that are applicable to that specific device type. In our example, we are adding OSA devices to ITSOTEST.

8. Complete the following updates to define the Processor / CU connection (see Figure 13-29), and press Enter:
   - `OFFLINE` to `No` (if you want these devices to be Online during IPL time).
   - `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   - `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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1            |
| Device number  . . : 1B80         Number of devices  : 15                   |
| 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 13-29 I/O Devices: Define Device Parameters / Features, OSD

9. The Assign/Unassign Device to Esoteric panel is now displayed where you can specify to which Esoteric (if any) you want the devices to be added. We add only the OSA/OSAD devices to the OSCONFIG ITSOTEST and not to any esoterics in this example (see Figure 13-30). Press Enter.

```
*----------- Define Device to Operating System Configuration -----------*
|                                                           Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : 1B80          Number of devices  : 15              |
| Device type  . . : OSA                                                |
|                                                                       |
| / Config. ID   Type     SS Description                     Defined    |
| _ COD230       MVS         z/OS image for COD test         Yes        |
| _ DBSV4SU4     MVS         z/OS DB Server 4                            |
| _ DBSV5SU4     MVS         z/OS DB Server 5                            |
| _ DBSV6SU4     MVS         z/OS DB Server 6                            |
| _ ITSO         MVS         All ITSO devices                Yes        |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes        |
| _ PERF4SU4     MVS         z/OS Appl Server 4                          |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                            |
| _ PERF6SU4     MVS         z/OS Appl Server 6                          |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh     Yes        |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                       |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                      |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image            Yes        |
*-----------------------------------------------------------------------*
```

Figure 13-30 I/O Devices: Assign/Unassign Device to Esoteric, OSD

10. The final panel is now displayed that shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

11. Perform the same action of adding OSAD device 1B8F to OSCONFIG ITSOTEST. Press Enter (see Figure 13-31).

```
*----------- Define Device to Operating System Configuration -----------*
|                                                             Row 1 of 42 |
| Command ===> ____________________________________ Scroll ===> CSR      |
|                                                                          |
| Select OSs to connect or disconnect devices, then press Enter.           |
|                                                                          |
| Device number  . : 1B8F          Number of devices  : 1                  |
| Device type  . . : OSAD                                                  |
|                                                                          |
| / Config. ID   Type     SS Description                      Defined      |
| _ COD230       MVS         z/OS image for COD test          Yes          |
| _ DBSV4SU4     MVS         z/OS DB Server 4                              |
| _ DBSV5SU4     MVS         z/OS DB Server 5                              |
| _ DBSV6SU4     MVS         z/OS DB Server 6                              |
| _ ITSO         MVS         All ITSO devices                 Yes          |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes          |
| _ PERF4SU4     MVS         z/OS Appl Server 4                            |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                              |
| _ PERF6SU4     MVS         z/OS Appl Server 6                            |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh      Yes          |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                         |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                        |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image             Yes          |
*-------------------------------------------------------------------------*
```

Figure 13-31 I/O Devices: Device / Processor Definition continued, OSD

13.2.8 Defining OSE CHPIDs

When defining an OSD or OSE connection, you must determine which type of OSA card you need for your configuration. OSD and ODE connections can use optical or copper-based cards, depending on the type and speed of network to which you want to connect your OSAs. The following cards are available:

- Optical based cards:
  - OSA Express6S GbE SX or LX
  - OSA Express6S 10 Gb SR or LR
- Copper-based card: OSA Express6S 1000BASE-T

A new OSE CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Partition access list
- For performance and redundancy, 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 (see the PCHID/CHID report for a list of installed hardware)

To define a new OSE 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 13-32) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 1                                                  |
|                                                                           |
| Channel path ID  . . . . B9    +           Channel ID  16C  +             |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSE   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . FC#0425 OSA Express6S 10GbE SR__                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 13-32 Processors: Add Channel Path, OSE

5. Complete the following updates and press Enter:
   - `Channel path ID` to `B9`.
   - `Channel ID` to `16C`.
   - `Channel path type` to `OSE`.
   - `Operational mode` to `SHR`.
   - `Description` to the description that you want.

6. HCD prompts you to select to which partition the CHPID should access. Enter a forward slash '/' next to the partition that you want (see Figure 13-33). 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 : 1                                                  |
| Channel path ID  . . : B9     Channel path type  . : OSE                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA1A          A      OS                                       |
| / 1      MUSCA1B          B      OS                                       |
| / 1      MUSCA1C          C      OS                                       |
| / 1      MUSCA1D          D      OS                                       |
| / 1      MUSCA1E          E      OS                                       |
| / 1      MUSCA1F          F      OS                                       |
| / 1      MUSCA11          1      OS                                       |
| / 1      MUSCA12          2      OS                                       |
| / 1      MUSCA13          3      OS                                       |
| / 1      MUSCA14          4      OS                                       |
| / 1      MUSCA15          5      OS                                       |
| / 1      MUSCA16          6      OS                                       |
*---------------------------------------------------------------------------*
```

Figure 13-33 Processors: Define Access List, OSE

7. Because more that one partition is defined in this CSS, HCD prompts us to define any Candidate list access. In our example, we so not Define any Candidate list access (see Figure 13-34 on page 321). Press Enter.

```
*-------------------------- Define Candidate List --------------------------*
|                                                              Row 1 of 15 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                          |
| Select one or more partitions for inclusion in the candidate list.       |
|                                                                          |
| Channel subsystem ID : 1                                                 |
| Channel path ID  . . : B9     Channel path type  . : OSE                 |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| _ 0      MUSCA0A          A      OS                                      |
| _ 0      MUSCA0B          B      OS                                      |
| _ 0      MUSCA0C          C      OS                                      |
| _ 0      MUSCA0D          D      OS                                      |
| _ 0      MUSCA0E          E      OS                                      |
| _ 0      MUSCA0F          F      OS                                      |
| _ 0      MUSCA01          1      OS                                      |
| _ 0      MUSCA02          2      OS                                      |
| _ 0      MUSCA03          3      OS                                      |
| _ 0      MUSCA04          4      OS                                      |
| _ 0      MUSCA05          5      OS                                      |
| _ 0      MUSCA06          6      OS                                      |
*--------------------------------------------------------------------------*
```

Figure 13-34 Processors: Define Access List, OSE

The HCD returns to the Channel Path List panel, which shows the CHPID that was just defined (see Figure 13-35).

```
 Goto  Filter  Backup  Query  Help
----------------------------------------------------------------------------
                              Channel Path List     Row 19 of 35 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 1

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ B3    12C   OSD   SPAN  __  __ __       No  1000BaseT
_ B4    170   OSD   SPAN  __  __ __       No  1000BaseT
_ B5    150   OSD   SPAN  __  __ __       No  1000BaseT
_ B8    148   OSD   SHR   __  __ __       No  FC#0425 OSA Express6s 10GbE SR
_ B9    16C   OSE   SHR   __  __ __       No  FC#0425 OSA Express6S 10GbE SR
_ BA    10C   OSD   SPAN  __  __ __       No  10GbE
_ BB    128   OSD   SPAN  __  __ __       No  10GbE
_ E0    12/1  CS5   SHR   __  __ __   Y   No  A09B-LG03/J.01
_ E1    12/2  CS5   SHR   __  __ __   Y   No  A09B-LG03/J.02
_ E4    16/1  CS5   SHR   __  __ __   Y   No  A09B-LG09/J.01
_ E5    16/2  CS5   SHR   __  __ __   Y   No  A09B-LG09/J.02
```

Figure 13-35 Processors: Channel Path List, OSE

13.2.9 Defining OSE CHPID connections to an OSA Control Unit

The only way to define an OSA connection to its Control Unit is direct connected.

You might want to connect the OSA Control Unit 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.

Consider the following points for connecting an OSE CHPID to an OSA Control Unit and its OSA/OSAD devices:

- How many OSAs are required to provide a primary and secondary/backup network connection
- Our example connects to a predefined OSA Control Unit (1B90) and OSA devices 1B90-1B9E and OSAD device 1B9F

To define OSE CHPID connections to an OSA Control Unit, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit you want to connect to, or in the command line enter `L 1B90`. In our example, we use `1B90`.
3. Enter `c` next to the Control Unit definition, and press Enter.
4. Make any changes to the Control Unit definition that you want (see Figure 13-36), and press Enter.

```
*----------------- Change Control Unit Definition ------------------*
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . 1B90  +                              |
|                                                                   |
| Control unit type  . . . . . OSA           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . 10G Copper - MUSCA                   |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
|                                                                   |
*-------------------------------------------------------------------*
```

Figure 13-36 Control units: Change Control Unit Definition, OSE

5. HCD now shows the Select Processor / CU panel, which is where the connection is now made between the OSE CHPID (B9) and the Control Unit (1B90). Enter `c` next to the Processor.CSS that contains the partition we want to have access to the control unit and also has access to the CHPID we want to connect to the control unit. Then, press **Enter**.

6. Update `Channel path IDs` to `B9` (channel path ID B9 is the CHPID that we defined in the previous example) to define the Processor / CU connection (see Figure 13-37). Press Enter.

```
*---------------------- Change Control Unit Definition -----------------------*
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Control unit number  . : 1B90         Type . . . . . . : OSA                  |
| Processor ID . . . . . : MUSCA        Musca                                  |
| Channel Subsystem ID . : 0                                                   |
|                                                                              |
| Channel path IDs . . . . B9    __   __   __   __   __   __   __     +        |
| 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 13-37 Control units: Change Control Unit Definition detail, OSE

7. HCD now shows the Modify Device Parameters panel where you can override the Unit Address (UA) numbers. For OSA definitions, the Unit address often starts at `00`.
8. Complete the following updates to define the Modify Device Parameters (see Figure 13-38), and press Enter:
   - `UA New` to `00` for the OSA device type.
   - `UA New` to `FE` for the OSAD device type.

```
*------------------------- Modify Device Parameters -------------------------*
|                                                  Row 1 of 2 More:       > |
| Command ===> __________________________________________ Scroll ===> CSR   |
|                                                                            |
| Specify or revise any changes to the device parameters in the list below. |
| To view attached control units, scroll to the right.                      |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                 |
| Channel Subsystem ID : 1                                                   |
|                                                                            |
| ---------Device---------     ---UA----                 Preferred Exposure  |
| No., Range Type          SS+ Old New + Time-Out STADET CHPID +   Device    |
| 1B90,015   OSA           _   00  00    No       No     __                  |
| 1B9F,001   OSAD          _   FE  FE    No       No     __                  |
| **************************** Bottom of data ****************************** |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 13-38 Control units: Modify Device Parameters, OSE

HCD returns to the Select Processor / CU panel, which shows the CHPID (B9) connection definition (see Figure 13-39).

```
Select Processor / CU    Row 1 of 12 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : 1B90     Control unit type . . . : OSA

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    B9      _______ _______ _______ _______ _______ _______ _______
_ MUSCA.2    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-39 Processors: Select Processor / CU: CHPID to Link address connection, OSE

9. By pressing F20 (Right), you can see the other parts of the definition summary.

13.2.10 Defining OSA and OSAD devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines which devices a z/OS system can access when it IPLs. The partition that the z/OS system is started in also must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups, which are defined in EDTs within an OSCONFIG. OSA definitions often do not use Esoterics.

Adding devices to an OSCONFIG includes the following considerations:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access that device.
- Our example adds the predefined OSA devices 1B90-1B9F (OSA/OSAD).

To define OSA and OSAD 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 you want to add to the OSCONFIG, or in the command line enter `L 1B90`. In our example, we use `1B90`.

   Because OSA and OSAD are two different device types, they must be added separately to the OSCONFIG.
3. Enter `c` next to the Device numbers, and press Enter.
4. HCD displays the Change Device Definition panel where you can modify the Control Unit to which the devices are attached (see Figure 13-40). Press Enter.

```
*------------------------- Change Device Definition --------------------------*
|                                                                              |
|                                                                              |
| Specify or revise the following values.                                      |
|                                                                              |
| Device number  . . . . . . . : 1B90  (0000 - FFFF)                           |
| Number of devices  . . . . . : 15                                            |
| Device type  . . . . . . . . : OSA                                           |
|                                                                              |
| Serial number  . . . . . . . . __________  +                                 |
| Description  . . . . . . . . . ________________________________              |
|                                                                              |
| Volume serial number . . . . . ______  + (for DASD)                          |
|                                                                              |
| PPRC usage . . . . . . . . . . _  + (for DASD)                               |
|                                                                              |
| Connected to CUs . 1B90  ____  ____  ____  ____  ____  ____  ____  +         |
|                                                                              |
| ENTER to continue.                                                           |
|                                                                              |
*------------------------------------------------------------------------------*
```

Figure 13-40 I/O Devices: Change Device Definition, OSE

5. HCD now displays the Device / Processor Definition panel where you can modify some of the Device parameters that relate to SubChannel Set, Unit Address, and Explicit Device Candidate List. Enter `c` next to the Proc.CCSID item if you want to modify any of these parameters (see Figure 13-41 on page 325) or press Enter to accept the default settings.

```
*---------------------- Device / Processor Definition -----------------------*
|                                                                 Row 1 of 1 |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                             |
| Select processors to change device/processor definitions, then press       |
| Enter.                                                                      |
|                                                                             |
| Device number  . . : 1B90       Number of devices  . : 15                  |
| Device type  . . . : OSA                                                    |
|                                                                             |
|                                      Preferred  Device Candidate List      |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit        Null  |
| _ MUSCA.1     _    00   No        No      __         No              ___   |
| **************************** Bottom of data ***************************** |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 13-41 I/O Devices: Device / Processor Definition continued, OSE

6. Next is the HCD panel where we Define Devices to the Operating System Configuration. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG you want to add to the Devices to, or in the command line enter `L ITSOTEST`. In our example, we use ITSOTEST.
7. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the Device Parameters and Features that are applicable to that specific device type. In our example, we are adding OSA devices to ITSOTEST.

8. Complete the following updates to define the Processor / CU connection (see Figure 13-42), and press Enter:
   - `OFFLINE` to `No` (if you want these devices to be Online during IPL time).
   - `DYNAMIC` to `Yes` (if you want the device to be changeable dynamically).
   - `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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1            |
| Device number  . . : 1B90         Number of devices  : 15                   |
| 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 13-42 I/O Devices: Define Device Parameters / Features, OSE

9. The Assign/Unassign Device to Esoteric panel is now displayed where you can specify to which Esoteric (if any) you want the devices to be added. We add only the OSA/OSAD devices to the OSCONFIG ITSOTEST and not to any esoterics in this example (see Figure 13-43). Press Enter.

```
*----------- Define Device to Operating System Configuration -----------*
|                                                             Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR     |
|                                                                         |
| Select OSs to connect or disconnect devices, then press Enter.         |
|                                                                         |
| Device number  . : 1B90          Number of devices  : 15               |
| Device type  . . : OSA                                                  |
|                                                                         |
| / Config. ID   Type     SS Description                      Defined    |
| _ COD230       MVS         z/OS image for COD test         Yes        |
| _ DBSV4SU4     MVS         z/OS DB Server 4                            |
| _ DBSV5SU4     MVS         z/OS DB Server 5                            |
| _ DBSV6SU4     MVS         z/OS DB Server 6                            |
| _ ITSO         MVS         All ITSO devices                Yes        |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes        |
| _ PERF4SU4     MVS         z/OS Appl Server 4                          |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                            |
| _ PERF6SU4     MVS         z/OS Appl Server 6                          |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh     Yes        |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                       |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                      |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image            Yes        |
*-------------------------------------------------------------------------*
```

Figure 13-43 I/O Devices: Assign/Unassign Device to Esoteric, OSE

10. The final panel is now displayed, which shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

11. Perform the same action of adding OSAD device 1B4F to OSCONFIG ITSOTEST.

12. Press Enter (see Figure 13-44).

```
*----------- Define Device to Operating System Configuration -----------*
|                                                            Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                        |
| Select OSs to connect or disconnect devices, then press Enter.         |
|                                                                        |
| Device number  . : 1B9F          Number of devices  : 1                |
| Device type  . . : OSAD                                                |
|                                                                        |
| / Config. ID   Type     SS Description                     Defined     |
| _ COD230       MVS         z/OS image for COD test         Yes         |
| _ DBSV4SU4     MVS         z/OS DB Server 4                             |
| _ DBSV5SU4     MVS         z/OS DB Server 5                             |
| _ DBSV6SU4     MVS         z/OS DB Server 6                             |
| _ ITSO         MVS         All ITSO devices                Yes         |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes         |
| _ PERF4SU4     MVS         z/OS Appl Server 4                           |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                             |
| _ PERF6SU4     MVS         z/OS Appl Server 6                           |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh     Yes         |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                        |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                       |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image            Yes         |
*------------------------------------------------------------------------*
```

Figure 13-44 I/O Devices: Device / Processor Definition continued, OSE

13.2.11 Defining OSM CHPIDs

When defining an OSM connection, you must determine which type of OSA card that you need for your configuration. OSM connections use the copper-based OSA Express6S 1000BASE-T card. The following cards are available:

- Optical based cards:
  - OSA Express6S GbE SX or LX
  - OSA Express6S 10 Gb SR or LR
- Copper-based card: OSA Express6S 1000BASE-T

The following are considerations for a new OSM CHPID:

- Channel path ID (CHPID)
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Partition access list
- For performance and redundancy, 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 (see the PCHID/CHID report for a list of installed hardware)

To define a new OSM 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 13-45) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 1                                                  |
|                                                                           |
| Channel path ID  . . . . B1    +            Channel ID  130  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSM   +                                          |
| Operation mode . . . . . SHR   +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . FC#0426 OSA Express6S 1000Base-T                 |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 13-45 Processors: Add Channel Path, OSM

5. Complete the following updates and press Enter:
   - `Channel path ID` to `B1`.
   - `Channel ID` to `130`.
   - `Channel path type` to `OSM`.
   - `Operational mode` to `SHR`.
   - `Description` to description that you want.
6. HCD now prompts you to specify whether the Allow more than 160 TCP/IP stacks. The default is `Yes`, which we use for our example (see Figure 13-46). 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?  . . . Yes                    |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 13-46 Processors: Allow for more than 160 TCP/IP stacks, OSM

7. HCD prompts you to select to which partition the CHPID should have access. Enter a forward slash '/' next to the partition that you want (see Figure 13-47). 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 : 1                                                  |
| Channel path ID  . . : B1     Channel path type  . : OSM                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA1A          A      OS                                       |
| / 1      MUSCA1B          B      OS                                       |
| / 1      MUSCA1C          C      OS                                       |
| / 1      MUSCA1D          D      OS                                       |
| / 1      MUSCA1E          E      OS                                       |
| / 1      MUSCA1F          F      OS                                       |
| / 1      MUSCA11          1      OS                                       |
| / 1      MUSCA12          2      OS                                       |
| / 1      MUSCA13          3      OS                                       |
| / 1      MUSCA14          4      OS                                       |
| / 1      MUSCA15          5      OS                                       |
| / 1      MUSCA16          6      OS                                       |
*---------------------------------------------------------------------------*
```

Figure 13-47 Processors: Define Access List, OSM

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.

8. HCD returns to the Channel Path List, which shows you the CHPID that was just defined (see Figure 13-48).

```
                              Channel Path List     Row 11 of 35 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 1

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 49    15D   FC    SPAN  02  02 1C       No  Disk *** was 02.24
_ 54    17D   FC    SPAN  01  01 0C       No  FCTC
_ 55    179   FC    SPAN  02  02 0C       No  FCTC
_ 58    13C   FC    SPAN  01  01 1B       No  Disk *** was 01.25
_ 59    13D   FC    SPAN  02  02 1D       No  Disk *** was 02.25
_ B0    110   OSC   SHR   __  __ __       No  FC#0426 OSA Express6S 1000Base-T
_ B1    130   OSM   SHR   __  __ __       No  FC#0426 OSA Express6S 1000Base-T
_ B2    14C   OSD   SPAN  __  __ __       No  1000BaseT
_ B3    12C   OSD   SPAN  __  __ __       No  1000BaseT
_ B4    170   OSD   SPAN  __  __ __       No  1000BaseT
_ B5    150   OSD   SPAN  __  __ __       No  1000BaseT
```

Figure 13-48 Processors: Channel Path List, OSM

13.2.12 Defining OSM CHPID connections to an OSM Control Unit

The only way to define an OSA connection to its Control Unit is direct connected.

You might want to connect the OSM Control Unit definition to multiple CPCs, even though the physical OSA is still unique to any one CPC. Also, you might want to span the OSM over multiple CSSs within a CPC.

Connecting an OSM CHPID to an OSM Control Unit and its OSA-M devices includes the following considerations:

- How many OSAs are required to provide a primary and secondary/backup network connection.
- Our example connects to a predefined OSM Control Unit (1B10) and OSA-M devices 1B10-1B1F.

To define OSM CHPID connections to an OSM Control Unit, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit 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 Control Unit definition, and press Enter.
4. Make any changes to the Control Unit definition that you want (see Figure 13-49 on page 332), and press Enter.

```
*----------------- Change Control Unit Definition ------------------*
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . 1B10  +                              |
|                                                                   |
| Control unit type  . . . . . OSM           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . ________________________________     |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
|                                                                   |
*-------------------------------------------------------------------*
```

Figure 13-49 Control units: Change Control Unit Definition, OSM

5. HCD now shows the Select Processor / CU panel. This panel is where the connection is now made between the OSD CHPID (B1) and the Control Unit (1B10). Enter `c` next to the Processor.CSS that contains the partition we want to have access to the control unit and also has access to the CHPID we want to connect to the control unit. Then, press Enter.
6. Update `Channel path IDs` to `B1` (channel path ID B1 is the CHPID we defined in the previous example) to define the Processor / CU connection (see Figure 13-50). Press Enter.

```
*---------------------- Change Control Unit Definition -----------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : 1B10         Type . . . . . . : OSM                |
| Processor ID . . . . . : MUSCA        Musca                                 |
| Channel Subsystem ID . : 1                                                  |
|                                                                             |
| Channel path IDs . . . . B1    __    __    __    __    __    __    __    +  |
| 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 13-50 Control units: Change Control Unit Definition detail, OSM

7. HCD now shows the Modify Device Parameters panel where you can override the UA numbers. For OSA definitions, the Unit address often starts at `00`. Update `UA New` to `00` for the OSA-M device type to define the Modify Device Parameters (see Figure 13-51). Press Enter.

```
*------------------------- Modify Device Parameters -------------------------*
|                                                   Row 1 of 1 More:       > |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Specify or revise any changes to the device parameters in the list below. |
| To view attached control units, scroll to the right.                       |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                  |
| Channel Subsystem ID : 1                                                   |
|                                                                            |
| ---------Device---------     ---UA----                 Preferred Exposure  |
| No., Range Type          SS+ Old New + Time-Out STADET CHPID +   Device     |
| 1B10,016   OSA-M         _   10  00    No       No     __                  |
| ***************************** Bottom of data ***************************** |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 13-51 Control units: Modify Device Parameters, OSM

HCD returns to the Select Processor / CU panel that shows the CHPID (B3) connection definition (see Figure 13-52).

```
                            Select Processor / CU    Row 1 of 12 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : 1B10     Control unit type . . . : OSM

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    B1      _______ _______ _______ _______ _______ _______ _______
_ MUSCA.2    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-52 Processors: Select Processor / CU: CHPID to Link address connection, OSM

8. Press F20 (Right) to see the other parts of the definition summary.

13.2.13 Defining OSA-M devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines which devices a z/OS system can access when it IPLs. The partition that the z/OS system is started in also must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups that are defined in EDTs within an OSCONFIG. OSA definitions often do not use Esoterics.

Adding devices to an OSCONFIG includes the following considerations:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access that device.
- Our example adds the predefined OSA devices 1B10-1B1F (OSA-M).

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 you want to add to the OSCONFIG or in the Command Line enter `L 1B10`. In our example, we use `1B10`.
3. Enter c next to the Device numbers, and press Enter.

HCD displays the Change Device Definition panel where you can modify the Control Unit to which the devices are attached (see Figure 13-53). Press Enter.

```
*------------------------ Change Device Definition -------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Device number  . . . . . . . : 1B10  (0000 - FFFF)                        |
| Number of devices  . . . . . : 16                                         |
| Device type  . . . . . . . . : OSA-M                                      |
|                                                                           |
| Serial number  . . . . . . . . __________  +                              |
| Description  . . . . . . . . . ________________________________           |
|                                                                           |
| Volume serial number . . . . . ______  + (for DASD)                       |
|                                                                           |
| PPRC usage . . . . . . . . . . _  + (for DASD)                            |
|                                                                           |
| Connected to CUs . 1B10  ____  ____  ____  ____  ____  ____  ____  +      |
|                                                                           |
| ENTER to continue.                                                        |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 13-53 I/O Devices: Change Device Definition, OSM

HCD now displays the Device / Processor Definition panel where you can modify some of the Device parameters that relate to SubChannel Set, Unit Address, and Explicit Device Candidate List.

Enter c next to the Proc.CCSID item if you want to modify any of these parameters (see Figure 13-54) 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  . . : 1B10       Number of devices  . : 16                |
| Device type  . . . : OSA-M                                               |
|                                                                          |
|                                          Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null  |
| _ MUSCA.1     _    00   No        No      __         No             ___   |
| ***************************** Bottom of data **************************** |
|                                                                          |
|                                                                          |
*---------------------------------------------------------------------------*
```

Figure 13-54 I/O Devices: Device / Processor Definition continued, OSM

Next is the HCD panel where we Define Devices to the Operating System Configuration.

4. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG you want to add to the Devices to or in the Command Line enter L ITSOTEST. In our example, we use ITSOTEST.
5. Enter s next to the OSCONFIG. Press Enter.

   HCD displays the Device Parameters and Features that are applicable to that specific device type. In our example, we are adding OSA-M devices to ITSOTEST.
6. Complete the following updates to define the Processor / CU connection (see Figure 13-55 on page 336), and press Enter:
   - OFFLINE to No (if you want these devices to be Online during IPL time).
   - DYNAMIC to Yes (if you want the device to be changeable dynamically).
   - 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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1            |
| Device number  . . : 1B10         Number of devices  : 16                   |
| Device type  . . . : OSA-M                                                  |
|                                                                             |
| 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 13-55 I/O Devices: Define Device Parameters / Features, OSM

The Assign/Unassign Device to Esoteric panel is now displayed where you can specify to which Esoteric (if any) you want the devices to be added.

We add only the OSA-M devices to the OSCONFIG ITSOTEST and not to any esoterics in this example.

7. Press Enter (see Figure 13-56).

   The final panel is now displayed showing that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

```
*----------- Define Device to Operating System Configuration -----------*
|                                                          Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR   |
|                                                                      |
| Select OSs to connect or disconnect devices, then press Enter.       |
|                                                                      |
| Device number  . : 1B10          Number of devices  : 16             |
| Device type  . . : OSA-M                                             |
|                                                                      |
| / Config. ID   Type     SS Description                     Defined   |
| _ COD230       MVS         z/OS image for COD test                   |
| _ DBSV4SU4     MVS         z/OS DB Server 4                          |
| _ DBSV5SU4     MVS         z/OS DB Server 5                          |
| _ DBSV6SU4     MVS         z/OS DB Server 6                          |
| _ ITSO         MVS         All ITSO devices                          |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes       |
| _ PERF4SU4     MVS         z/OS Appl Server 4                        |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                          |
| _ PERF6SU4     MVS         z/OS Appl Server 6                        |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh               |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                     |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                    |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image                      |
*----------------------------------------------------------------------*
```

Figure 13-56 I/O Devices: Assign/Unassign Device to Esoteric, OSM

13.3 IQD CHPIDs for Hypersockets

This section describes the process to define IQD CHPIDs and their control units and devices.

13.3.1 Defining IQD CHPIDs

When Hypersockets must be defined, the CHPID type of IQD (Internal Queued Direct Communication) is used. Since z13, the IQD CHPID also requires the VCHID statement. The valid range for the VCHIDs are 7E0 - 7FF.

As with ICP CHPIDs, IQD CHPIDs are defined logically and internally to the processor and require no installed hardware (PCHIDs). However, from z13, a maximum of up to 32 high-speed virtual LAN attachments can be defined.

Each of the logical IQD VCHIDs can support only 1 CHPID; however, the CHPIDs can be spanned across multiple CSSs.

A new IQD CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel ID (vCHID)
- Channel path type
- Operational mode
- Description
- Partition access list

To define a new IQD 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 13-57 on page 338) to add a new CHPID.
5. Complete the following updates and press Enter:
   - `Channel path ID` to `F1`.
   - `Channel ID` to `7C1`.
   - `Channel path type` to `IQD`.
   - `Operational mode` to `SHR`.
   - `Description` to the description that you want.

```
*--------------------- Change Channel Path Definition ----------------------*
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID . . . . : MUSCA       Musca                                   |
| Configuration mode . : LPAR                                                |
| Channel Subsystem ID : 1                                                   |
|                                                                            |
| Channel path ID  . . . . F1    +           Channel ID  7C1  +              |
| 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 13-57 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 Hypersockets** for the IQD function (see Figure 13-58). Press Enter.

```
*------------------- Specify IQD Channel Parameters ---------------------*
|                                                                        |
|                                                                        |
| Specify or revise the values below.                                    |
|                                                                        |
| Maximum frame size in KB . . . . . . 64  +                             |
|                                                                        |
| IQD function . . . . . . . . . . . . 1   1.  Basic HiperSockets        |
|                                          2.  IEDN Access (IQDX)        |
|                                          3.  External Bridge           |
|                                                                        |
| Physical network ID  . . . . . . . . ________________                  |
|                                                                        |
|                                                                        |
*------------------------------------------------------------------------*
```

Figure 13-58 Processors: Specify IQD Channel Parameters, IQD

HCD prompts you to select to which partitions the CHPID can access.

7. Enter a forward slash '/' next to the wanted partitions (see Figure 13-59), 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 : 1                                                  |
| Channel path ID  . . : F1     Channel path type  . : IQD                  |
| Operation mode . . . : SHR    Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA1A          A      OS                                       |
| / 1      MUSCA1B          B      OS                                       |
| / 1      MUSCA1C          C      OS                                       |
| / 1      MUSCA1D          D      OS                                       |
| _ 1      MUSCA1E          E      OS                                       |
| _ 1      MUSCA1F          F      OS                                       |
| / 1      MUSCA11          1      OS                                       |
| / 1      MUSCA12          2      OS                                       |
| / 1      MUSCA13          3      OS                                       |
| / 1      MUSCA14          4      OS                                       |
| / 1      MUSCA15          5      OS                                       |
| / 1      MUSCA16          6      OS                                       |
*---------------------------------------------------------------------------*
```

Figure 13-59 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.

HCD now returns to the Channel Path List that shows you the CHPID that was just defined (see Figure 13-60).

```
                              Channel Path List     Row 35 of 35 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 1

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ F1    7C1   IQD   SHR   __  __ __       No  IQD Internal Queued Direct Comms
```

Figure 13-60 Processors: Channel Path List, IQD

13.3.2 Defining IQD CHPID connections to an IQD Control Unit

The only way to define an IQD connection to its Control Unit is direct connected.

You might want to span the IQD CHPID over multiple CSSs within a CPC.

Connecting an IQD CHPID to an IQD Control Unit and its IQD devices includes the following considerations:

- How many IQDs are required to provide the required Hypersocket bandwidth.
- Our example connects to a predefined IQD Control Unit (F100) and IQD devices 0FC0-0FCF.

To define IQD CHPID connections to an IQD Control Unit, complete the following steps:

1. From the main HCD panel, select option **1.4. Control units**.
2. Scroll through the Control Unit List until you find the control unit you want to connect to, or in the command line enter `L F100`. In our example, we use `F100`.
3. Enter `c` next to the Control Unit definition, and press Enter.
4. Make any changes to the Control Unit definition that you want (see Figure 13-61), and press Enter.

```
*----------------- Change Control Unit Definition ------------------*
|                                                                   |
|                                                                   |
| Specify or revise the following values.                           |
|                                                                   |
| Control unit number  . . . . F100  +                              |
|                                                                   |
| Control unit type  . . . . . IQD           +                      |
|                                                                   |
| Serial number  . . . . . . . __________  +                        |
| Description  . . . . . . . . ________________________________     |
|                                                                   |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +           |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +           |
|                                                                   |
| Define more than eight ports . . 2   1.  Yes                      |
|                                      2.  No                       |
|                                                                   |
*-------------------------------------------------------------------*
```

Figure 13-61 Control units: Change Control Unit Definition, IQD

5. HCD now shows the Select Processor / CU panel. This panel is where the connection is now made between the IQD CHPID (F1) and the Control Unit (F100). Enter `c` next to the Processor.CSS that contains the partition you want to have access to the control unit and also has access to the CHPID you want to connect to the control unit. Then, press Enter.

6. Update `Channel path IDs` to `E6` (channel path ID E6 is the CHPID we defined in the previous example) to define the Processor / CU connection (see Figure 13-62). Press Enter.

```
*---------------------- Change Control Unit Definition -----------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : F100         Type . . . . . . : IQD                 |
| Processor ID . . . . . : MUSCA        Musca                                 |
| Channel Subsystem ID . : 1                                                  |
|                                                                             |
| Channel path IDs . . . . F1    __    __    __    __    __    __    __     + |
| 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 13-62 Control units: Change Control Unit Definition detail, IQD

7. HCD now shows the Modify Device Parameters panel where you can override the UA numbers. For IQD definitions, the Unit address often starts at `00`.

   Update `UA New` to `00` for the IQD device type to define the Modify Device Parameters (see Figure 13-63). Press Enter.

```
*------------------------- Modify Device Parameters --------------------------*
|                                                  Row 1 of 1 More:       >  |
| Command ===> __________________________________________ Scroll ===> CSR    |
|                                                                            |
| Specify or revise any changes to the device parameters in the list below.  |
| To view attached control units, scroll to the right.                       |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                  |
| Channel Subsystem ID : 1                                                   |
|                                                                            |
| ---------Device---------     ---UA----                 Preferred Exposure  |
| No., Range Type          SS+ Old New + Time-Out STADET CHPID +   Device    |
| 0FC0,032   IQD           _   C0  00    No       No     __                  |
| ***************************** Bottom of data ***************************** |
|                                                                            |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 13-63 Control units: Modify Device Parameters, IQD

HCD now returns to the Select Processor / CU panel that shows the CHPID (E6) connection definition (see Figure 13-64).

```
 Select Processor / CU    Row 1 of 12 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Control unit number . . : F100     Control unit type . . . : IQD

             ---------------Channel Path ID . Link Address + ---------------
/ Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
_ CETUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.3    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.4    _______ _______ _______ _______ _______ _______ _______ _______
_ CETUS.5    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.0    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.1    _______ _______ _______ _______ _______ _______ _______ _______
_ LEPUS.2    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.0    _______ _______ _______ _______ _______ _______ _______ _______
_ MUSCA.1    F1      _______ _______ _______ _______ _______ _______ _______
_ MUSCA.2    _______ _______ _______ _______ _______ _______ _______ _______
```

Figure 13-64 Processors: Select Processor / CU: CHPID to Link address connection, IQD

8. Press F20 (Right) to see the other parts of the definition summary.

13.3.3 Defining IQD devices to an OSCONFIG

The OSCONFIG name is the part of an IODF that determines which devices a z/OS system can access when it IPLs. The partition that the z/OS system is started in also must access the CHPIDs that connect to the Control Units and Devices that match in the OSCONFIG.

The OSCONFIG also contains Esoterics device groups that are defined in EDTs within an OSCONFIG. IQD definitions often do not use Esoterics.

Adding devices to an OSCONFIG includes the following considerations:

- Adding a device to an OSCONFIG does not necessarily mean that the z/OS system can access that device.
- Our example adds the predefined IQD devices 0FC0-0FDF (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 F100`. In our example, we use `F100`.
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 Control Unit to which the devices are attached (see Figure 13-65 on page 343). Press Enter.

```
*------------------------- Change Device Definition --------------------------*
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number  . . . . . . . : 0FC0  (0000 - FFFF)                          |
| Number of devices  . . . . . : 32                                           |
| Device type  . . . . . . . . : IQD                                          |
|                                                                             |
| Serial number  . . . . . . . . __________  +                                |
| Description  . . . . . . . . . ________________________________             |
|                                                                             |
| Volume serial number . . . . . ______  + (for DASD)                         |
|                                                                             |
| PPRC usage . . . . . . . . . . _  + (for DASD)                              |
|                                                                             |
| Connected to CUs . F100  ____  ____  ____  ____  ____  ____  ____  +        |
|                                                                             |
| ENTER to continue.                                                          |
|                                                                             |
|                                                                             |
*-----------------------------------------------------------------------------*
```

Figure 13-65 I/O Devices: Change Device Definition, IQD

5. HCD now displays the Device / Processor Definition panel where you can modify some of the Device parameters that relate to SubChannel Set, Unit Address, and Explicit Device Candidate List. Enter `c` next to the Proc.CCSID item if you want to modify any of these parameters (see Figure 13-66), 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  . . : 0FC0       Number of devices  . : 32                 |
| Device type  . . . : IQD                                                  |
|                                                                           |
|                                          Preferred  Device Candidate List |
| / Proc.CSSID  SS+  UA+  Time-Out  STADET  CHPID +    Explicit       Null  |
| _ MUSCA.1     _    00   No        No      __         No             ___   |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 13-66 I/O Devices: Device / Processor Definition continued, IQD

6. The HCD panel where we Define Devices to the Operating System Configuration is displayed. Scroll through the list of OSCONFIG definitions until you find the OSCONFIG to which you want to add to the Devices, or in the command line enter `L ITSOTEST`. In our example, we use ITSOTEST.
7. Enter `s` next to the OSCONFIG, and press Enter.

   HCD displays the Device Parameters and Features that are applicable to that specific device type. In our example, we are adding IQD devices to ITSOTEST.
8. Complete the following updates to define the Processor / CU connection (see Figure 13-67), and press Enter:
   - `OFFLINE` to No (if you want these devices to be Online during IPL time).
   - `DYNAMIC` to Yes (if you want the device to be changeable dynamically).
   - `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 . : ITSOTEST     ITSO test OSCONFIG for z14 ZR1           |
| Device number  . . : 0FC0         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 13-67 I/O Devices: Define Device Parameters / Features, IQD

9. The Assign/Unassign Device to Esoteric panel is now displayed where you can specify to which Esoteric (if any) you want the devices to be added. We add only the IQD devices to the OSCONFIG ITSOTEST and not to any esoterics in this example. Press Enter (see Figure 13-68).

```
*----------- Define Device to Operating System Configuration -----------*
|                                                              Row 1 of 42 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                        |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                        |
| Device number  . : 0FC0          Number of devices  : 32              |
| Device type  . . : IQD                                                |
|                                                                        |
| / Config. ID   Type     SS Description                      Defined   |
| _ COD230       MVS         z/OS image for COD test                      |
| _ DBSV4SU4     MVS         z/OS DB Server 4                             |
| _ DBSV5SU4     MVS         z/OS DB Server 5                             |
| _ DBSV6SU4     MVS         z/OS DB Server 6                             |
| _ ITSO         MVS         All ITSO devices                             |
| _ ITSOTEST     MVS         ITSO test OSCONFIG for z14 ZR1   Yes         |
| _ PERF4SU4     MVS         z/OS Appl Server 4                           |
| _ PERF5SU4     MVS         z/OS Perf Serv 5                             |
| _ PERF6SU4     MVS         z/OS Appl Server 6                           |
| _ ZMEDRVR      MVS         z/OS 2.1 Server Pak refresh                  |
| _ ZOSADLTD     MVS         z/OS 1.13 ADLT system                        |
| _ ZOSALDTE     MVS         z/OS 2.1 ADLT-E system                       |
| _ ZOSDBSV1     MVS         z/OS 2.1 dbsv1 Image                         |
*------------------------------------------------------------------------*
```

Figure 13-68 I/O Devices: Assign/Unassign Device to Esoteric, IQD

10. The final panel is now displayed, which shows that the devices are defined to the OSCONFIG. Press Enter to return to the I/O Device List.

14

Adding coupling connectivity

This chapter describes the steps to define CS5, CL5, and ICP CHPIDs, and Coupling Facility links between these CHPID types.

It also provides a list of these potential configuration items and a short description of how to do each of them by using hardware configuration definition (HCD).

Note: Not all of the configuration items that are described in this chapter are necessarily required for your installation. Also, the examples that are presented are not intended to be exhaustive.

This chapter includes the following topics:

14.1 Defining more I/O using HCD

When defining new I/O components in an IODF, certain definitions, such as OSCONFIGs, partitions, FICON Switches, Control Units, and Devices must be made up front. After the following definitions are made, the connections can be made:

- Defining Coupling Facility LPARs in a CSS
- Defining CS5 CHPIDs
- Defining a Coupling Facility link with CS5 CHPIDs
- Defining CL5 CHPIDs
- Defining a Coupling Facility link with CL5 CHPIDs
- Defining ICP CHPIDs
- Defining a Coupling Facility link with ICP CHPIDs

The I/O definitions that are described next use HCD to demonstrate the examples. The examples continue by using the example that was created in Chapter 5, “Production IODF and setting up the CPC” on page 85 (for example, `ITSO.IODF78.WORK`).

14.2 Coupling Facility LPARs, CS5, CL5, and ICP CHPIDs

This section covers defining Coupling Facility LPARs and the CS5, CL5, and ICP CHPID type definitions.

14.2.1 Defining Coupling Facility LPARs in a CSS

A new (unreserved) partition includes the following considerations:

- Partition name
- Number
- Usage
- Description
- To add CHPIDs to a partition, they first must be defined to the processor
- Renaming a partition is a two-step process:
  a. It must be redefined as reserved (Partition name = *) and the IODF activated on the processor.
  b. The partition must be redefined to the new name and the IODF activated on the processor.

To change a reserved partition to an active partition in a 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 A for this example), and press Enter.

5. Complete the following updates (see Figure 14-1) and press Enter:
   – Update `Partition Name` to `MUSCA2A` (a naming standard based on CSS=2, Partition =A).
   – Review `Partition usage` and change if required. We use CF in this example.
   – Update `Description` to `MUSCA2A test CF partition`.

```
*---------------------------- 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 |                                                          |  |
| c MUSCA2A | Partition name . . . MUSCA2A                             |  |
| _ *       | Partition number . . A      (same as MIF image ID)       |  |
| _ *       | Partition usage  . . CF     +                            |  |
| _ *       | UID uniqueness . . . N        (Y/N)                      |  |
| _ *       |                                                          |  |
| _ *       | Description  . . . . MUSCA2A test CF partition           |  |
| _ *       |                                                          |  |
| _ *       |                                                          |  |
| _ *       |                                                          |  |
| _ *       |                                                          |  |
*---------- *----------------------------------------------------------* --*
```

Figure 14-1 Processors: Change Partition Definition, Coupling Facility

14.2.2 Defining CS5 CHPIDs

When defining a CS5 CHPID to create a Coupling Facility link between a Coupling Facility LPAR and a z/OS LPAR, first determine which z/OS LPARs require access to which CF LPAR, how many CF links are required, and to how many different physical processors.

CS5 coupling facility CHPIDs are defined by using FC 0172 (ICA SR 2 Links) cards that are installed on the CPC drawer as opposed to 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.

A new CS5 CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel path type
- Operational mode
- Description
- Adapter ID (AID)
- Adapter Port
- Partition access list

- For performance and redundancy, how many AID cards are installed in the processor and to what PCIe slot on what CPC drawer (see the PCHID/CHID report for a list of installed hardware)

To define a new 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. In the command line, enter `add` (see Figure 14-2) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 2                                                  |
|                                                                           |
| Channel path ID  . . . . E0    +           Channel ID ___  +              |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . CS5   +                                          |
| Operation mode . . . . . SPAN  +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . FC#0172 ICA SR 2 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

5. Complete the following updates and press Enter:
   - `Channel path ID` to `E0`
   - `Channel path type` to `CS5`
   - `Operational mode` to `SPAN`
   - `Description` to the description that you want

6. HCD prompts you to specify the adapter and port of the HCA attributes. Complete the following updates (see Figure 14-3) and press Enter:
   - Adapter of the HCA to 12
   - Port on the HCA to 1

```
*------------------ Specify HCA Attributes -------------------*
|                                                             |
|                                                             |
| Specify or revise the values below.                         |
|                                                             |
| Adapter ID of the HCA  . . 12  +                            |
| Port on the HCA  . . . . . 1   +                            |
|                                                             |
|                                                             |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 14-3 Processors: Specify HCA Attributes, CS5

7. HCD prompts you to select to which partitions the CHPID should have access. Enter a forward slash '/' next to the wanted partitions (see Figure 14-4), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                              Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.          |
|                                                                           |
| Channel subsystem ID : 2                                                  |
| Channel path ID  . . : E0     Channel path type  . : CS5                  |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA11          1      OS                                       |
| _ 1      MUSCA12          2      OS                                       |
| _ 1      MUSCA13          3      OS                                       |
| _ 1      MUSCA14          4      OS                                       |
| _ 1      MUSCA15          5      OS                                       |
| _ 1      MUSCA16          6      OS                                       |
| _ 1      MUSCA17          7      OS                                       |
| _ 1      MUSCA18          8      OS                                       |
| _ 1      MUSCA19          9      OS                                       |
| / 2      MUSCA2A          A      CF    MUSCA2A test CF partition         |
| **************************** Bottom of data ***************************** |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 14-4 Processors: Define Access List, CS5

8. Because more partitions than the selected two are defined, the next panel `Define Candidate List` appears. For this example, we do not add any partitions in the access list. Therefore, press Enter. HCD now returns to the Channel Path List, which shows you the CHPID that was defined (see Figure 14-5).

```
                              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 . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   N   No  FC#0172 ICA SR 2 Links
```

Figure 14-5 Processors: Channel Path List, CS5

9. Define another CS5 CHPID as E4 to the same LPARs as AID=16, Port=1 (see Figure 14-6).

```
                              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 . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   N   No  FC#0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   N   No  FC#0172 ICA SR 2 Links
```

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 connected.

Creating a CF link by using CS5 CHPIDs includes the following considerations:

- The ICA SR connection is a physical cable between two FC 0172 cards on the same or different processors.
- Up to four logical CHPIDs per port can be defined over that physical connection.
- How many CS5 CF links are required to provide enough primary and secondary links and coupling bandwidth.
- Which z/OS LPARs on the same or different processors need access to the CF LPARs.
- CF links also provide Server Time Protocol (STP) connectivity for a Coordinated Time Network (CTN) between processors and a sysplex.
- Our example connects two CS5 CHPIDs (E0 and E4) on the same processor.

To define a Coupling Facility 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 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 you want to connect from, or in the command line, enter `L E0`. In our example, we use `E0`.
5. Enter `f` next to the CHPID definition (see Figure 14-7), and press Enter.

```
                              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 . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+           Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f E0    12/1  CS5   SPAN  __  __ __   N   No  FC#0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   N   No  FC#0172 ICA SR 2 Links
```

Figure 14-7 Processors: CF Channel Path Connectivity List, CS5

6. In the next panel, HCD prompts you to specify the second CS5 CHPID to which you want to connect. Scroll through the Channel Path List until you find the second CS5 CHPID you want to connect to, or in the command line enter `L E4`. In our example, we use `E4`.
7. Enter `p` next to the CHPID definition (see Figure 14-8), and press Enter.

```
                     CF Channel Path Connectivity List             Row 1 of
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ E0  12/1  Y  CS5  SPAN N
p E4  16/1  Y  CS5  SPAN N
```

Figure 14-8 Processors: CF Channel Path Connectivity List, CS5

8. HCD prompts you to specify to which CF channel path to connect. Complete the following updates (see Figure 14-9) and press Enter:
   - `Destination processor ID` to MUSCA
   - `Destination channel subsystem ID` to `2`
   - `Destination channel path ID` to `E0`

```
*---------------- Connect to CF Channel Path -----------------*
|                                                             |
|                                                             |
| Specify the following values.                               |
|                                                             |
| Source processor ID  . . . . . : MUSCA                      |
| Source channel subsystem ID  . : 2                          |
| Source channel path ID . . . . : E4                         |
| Source channel path type . . . : CS5                        |
|                                                             |
| Destination processor ID . . . . . . MUSCA___  +            |
| Destination channel subsystem ID . . 2   +                  |
| Destination channel path ID  . . . . E0  +                  |
|                                                             |
| Timing-only link . . . . . . . . . . No                     |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 14-9   Processors: Connect to CF Channel Path, CS5

9. HCD checks available Control Unit 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 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 . . . . . . . : MUSCA                                 |
| Channel subsystem ID . . . : 2                                     |
| Channel path ID  . . . . . : E4        Operation mode . . : SPAN   |
| Channel path type  . . . . : CS5                                   |
|                                                                    |
| Control unit number  . . . . FFFD  +                               |
|                                                                    |
| Device number  . . . . . . . FFDC                                  |
| Number of devices  . . . . . 8                                     |
|                                                                    |
|                                                                    |
*--------------------------------------------------------------------*
```

Figure 14-10   Processors: Add CF Control Unit and Devices, CS5

10. HCD provides suggested Control Unit numbers and Device addresses for the first CS5 CHPID (see Figure 14-11). The Control unit 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 . . . . . . . : MUSCA                                  |
| Channel subsystem ID . . . : 2                                      |
| Channel path ID  . . . . . : E0        Operation mode . . : SPAN   |
| Channel path type  . . . . : CS5                                    |
|                                                                      |
| Control unit number  . . . . FFFD  +                                |
|                                                                      |
| Device number  . . . . . . . FFD4                                   |
| 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
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID   CF Type Mode Occ   Proc.CSSID CHP CHID   CF Type Mode   Type  Dev
_ E0  12/1   Y  CS5  SPAN N     MUSCA.2    E4  16/1   Y  CS5  SPAN   CFP   8
_ E4  16/1   Y  CS5  SPAN N     MUSCA.2    E0  12/1   Y  CS5  SPAN   CFP   8
```

Figure 14-12 Processors: Add CF Channel Path Connectivity List, CS5

14.2.4 Defining CL5 CHPIDs

When defining a CL5 CHPID to create a Coupling Facility link between a Coupling Facility LPAR and a z/OS LPAR, determine which z/OS LPARs require access to which CF LPAR, how many CF links are required, and to how many different physical processors.

CL5 coupling facility CHPIDs are defined by using FC 0433 (Coupling Express LR) cards, which are installed in the PCIe+ I/O drawer instead of on the CPC drawer.

The Coupling Express LR card has two ports (Port 1 and Port 2), which provide two physical connections between another Coupling Express LR card on the same or different processor.

Each of the ports can have up to four CHPIDs defined.

A new CL5 CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel ID (CHID)
- Channel path type
- Operational mode
- Description
- Partition access list
- For performance and redundancy, how many CL5 cards are installed in the processor and to what PCIe slot on what CPC drawer (see the PCHID/CHID report for a list of installed hardware)

To define a new 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 command line, enter `add` (see Figure 14-13) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID . . . . : MUSCA        Musca                                  |
| Configuration mode . : LPAR                                                |
| Channel Subsystem ID : 2                                                   |
|                                                                            |
| Channel path ID  . . . . E8    +            Channel ID  ___  +             |
| Number of CHPIDs . . . . 1                                                 |
| Channel path type  . . . CL5   +                                           |
| Operation mode . . . . . SPAN  +                                           |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +        |
| Description  . . . . . . FC#0433 Coupling Express LR ____                  |
|                                                                            |
| 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

5. Complete the following updates and press Enter:
   - `Channel path ID` to `E8`
   - `Channel ID` to `124`
   - `Channel path type` to `CL5`
   - `Operational mode` to `SPAN`
   - `Description` to the description that you want

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  . . . 124                              |
| Coupling port  . . . . . . . 1   +                          |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 14-14 Processors: Specify Coupling Pchid/Port Attributes, CL5

7. HCD prompts you to select to which partitions the CHPID can access. Enter a forward slash '/' next to the partitions that you want (see Figure 14-15), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                            Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                          |
| Select one or more partitions for inclusion in the access list.          |
|                                                                          |
| Channel subsystem ID : 2                                                 |
| Channel path ID  . . : E8     Channel path type  . : CL5                 |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                   |
|                                                                          |
| / CSS ID Partition Name   Number Usage Description                       |
| / 1      MUSCA11          1      OS                                      |
| _ 1      MUSCA12          2      OS                                      |
| _ 1      MUSCA13          3      OS                                      |
| _ 1      MUSCA14          4      OS                                      |
| _ 1      MUSCA15          5      OS                                      |
| _ 1      MUSCA16          6      OS                                      |
| _ 1      MUSCA17          7      OS                                      |
| _ 1      MUSCA18          8      OS                                      |
| _ 1      MUSCA19          9      OS                                      |
| / 2      MUSCA2A          A      CF    MUSCA2A test CF partition         |
| **************************** Bottom of data **************************** |
|                                                                          |
*--------------------------------------------------------------------------*
```

Figure 14-15 Processors: Define Access List, CL5

8. Because more partitions than the selected two are defined, the next panel `Define Candidate List` appears. For this example, we do not add any partitions in the access list. Therefore, press Enter. HCD returns to the Channel Path List, which shows you the CHPID that was defined (see Figure 14-16 on page 358).

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

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   Y   No  FC#0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E8    124/1 CL5   SPAN  __  __ __   N   No  FC 0433 Coupling Express LR
```

Figure 14-16 Processors: Channel Path List, CL5

9. Define another CL5 CHPID as ED to the same LPARs as CHID=144, Port=1 (see Figure 14-17).

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

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E8    124/1 CL5   SPAN  __  __ __   N   No  FC 0433 Coupling Express LR
_ ED    144/2 CL5   SPAN  __  __ __   N   No  FC 0433 Coupling Express LR
```

Figure 14-17 Processors: Channel Path List, CL5

14.2.5 Defining a Coupling Facility link with CL5 CHPIDs

The only way to define a CL5 CHPID to another CL5 CHPID is direct connected.

Creating a CF link by using CL5 CHPIDs includes the following considerations:

- The Coupling Express LR connection is a physical cable between two FC 0433 cards on the same or different processors.
- Up to four logical CHPIDs per port can be defined over that physical connection.
- How many CL5 CF links are required to provide enough primary and secondary links and coupling bandwidth.
- 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.
- Our example connects two CL5 CHPIDs (E8 and ED) on the same processor.

To define a Coupling Facility link with CL5 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 includes the CL5 CHPID definition for which you want to create the first CF link, and press Enter.
4. Scroll through the Channel Path List until you find the first CL5 CHPID you want to connect from, or in the command line, enter `L E8`. In our example, we use `E8`.
5. Enter `f` next to the CHPID definition (see Figure 14-18) and press Enter.

```
                              Channel Path List       Row 3 of 4 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f E8    124/1 CL5   SPAN  __  __ __   N   No  FC 0433 Coupling Express LR
_ ED    144/2 CL5   SPAN  __  __ __   N   No  FC 0433 Coupling Express LR
```

Figure 14-18 Processors: CF Channel Path Connectivity List, CL5

6. HCD prompts you to specify the second CL5 CHPID to which you want to connect. Scroll through the Channel Path List until you find the second CL5 CHPID you want to connect to, or in the command line enter `L ED`. In our example, we use `ED`.
7. Enter `p` next to the CHPID definition (see Figure 14-19) and press Enter.

```
                     CF Channel Path Connectivity List             Row 3 of
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ E8  124/1 Y  CL5  SPAN N
p ED  144/2 Y  CL5  SPAN N
```

Figure 14-19 Processors: CF Channel Path Connectivity List, CL5

8. HCD prompts you to specify to which CF channel path to connect. Complete the following updates (see Figure 14-20) and press Enter:
   - `Destination processor ID` to `MUSCA`
   - `Destination channel subsystem ID` to `2`
   - `Destination channel path ID` to `E8`

```
*---------------- Connect to CF Channel Path -----------------*
|                                                             |
|                                                             |
| Specify the following values.                               |
|                                                             |
| Source processor ID  . . . . . : MUSCA                      |
| Source channel subsystem ID  . : 2                          |
| Source channel path ID . . . . : ED                         |
| Source channel path type . . . : CL5                        |
|                                                             |
| Destination processor ID . . . . . . MUSCA___  +            |
| Destination channel subsystem ID . . 2   +                  |
| Destination channel path ID  . . . . E8  +                  |
|                                                             |
| Timing-only link . . . . . . . . . . No                     |
|                                                             |
|                                                             |
*-------------------------------------------------------------*
```

Figure 14-20 Processors: Connect to CF Channel Path, CL5

9. HCD checks available Control Unit 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 14-21) 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 . . . . . . . : MUSCA                                 |
| Channel subsystem ID . . . : 2                                     |
| Channel path ID  . . . . . : ED         Operation mode . . : SPAN  |
| Channel path type  . . . . : CL5                                   |
|                                                                    |
| Control unit number  . . . . FFFD  +                               |
|                                                                    |
| Device number  . . . . . . . FFCC                                  |
| Number of devices  . . . . . 8                                     |
|                                                                    |
|                                                                    |
*--------------------------------------------------------------------*
```

Figure 14-21 Processors: Add CF Control Unit and Devices, CL5

10. HCD provides suggested Control Unit numbers and Device addresses for the first CL5 CHPID (see Figure 14-22). The Control unit 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 . . . . . . . : MUSCA                                  |
| Channel subsystem ID . . . : 2                                      |
| Channel path ID  . . . . . : E8        Operation mode . . : SPAN    |
| Channel path type  . . . . : CL5                                    |
|                                                                     |
| Control unit number  . . . . FFFD  +                                |
|                                                                     |
| Device number  . . . . . . . FFC4                                   |
| Number of devices  . . . . . 8                                      |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-22 Processors: Add CF Control Unit and Devices, CL5

HCD provides a summary of the Source and Destination CF Connectivity List (see Figure 14-23).

```
                     CF Channel Path Connectivity List             Row 3 of 4
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ E8  124/1 Y  CL5  SPAN N     MUSCA.2    ED  144/2 Y  CL5  SPAN   CFP   8
_ ED  144/2 Y  CL5  SPAN N     MUSCA.2    E8  124/1 Y  CL5  SPAN   CFP   8
```

Figure 14-23 Processors: Add CF Channel Path Connectivity List, CL5

14.2.6 Defining ICP CHPIDs

When defining an ICP CHPID to create a Coupling Facility link between a Coupling Facility LPAR and a z/OS LPAR, determine which z/OS LPARs require access to which CF LPAR, and how many CF links are required within the same physical processor.

ICP coupling facility CHPIDs are defined logically and internally to the processor and require no installed hardware. However, the maximum number of ICP CHPIDS that can be defined in a processor is limited to 32.

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.

A new ICP CHPID includes the following considerations:

- Channel path ID (CHPID)
- Channel path type
- Operational mode
- Description
- Partition access list

To define a new 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 command line, enter `add` (see Figure 14-24) to add a CHPID.

```
*---------------------------- Add Channel Path -----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : MUSCA        Musca                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 2                                                  |
|                                                                           |
| Channel path ID  . . . . EE    +            Channel ID  ___  +            |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . ICP   +                                          |
| Operation mode . . . . . SPAN  +                                          |
| 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 14-24 Processors: Add Channel Path, ICP

5. Complete the following updates and press Enter:
   - `Channel path ID` to `EE`
   - `Channel path type` to `ICP`
   - `Operational mode` to `SPAN`
   - `Description` to the description that you want

6. HCD prompts you to select to which partitions the CHPID should have access. Enter a forward slash '/' next to the partitions that you want (see Figure 14-25), and press Enter.

```
*--------------------------- Define Access List ----------------------------*
|                                                              Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.           |
|                                                                           |
| Channel subsystem ID : 2                                                  |
| Channel path ID  . . : EE     Channel path type  . : ICP                  |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA11          1      OS                                       |
| _ 1      MUSCA12          2      OS                                       |
| _ 1      MUSCA13          3      OS                                       |
| _ 1      MUSCA14          4      OS                                       |
| _ 1      MUSCA15          5      OS                                       |
| _ 1      MUSCA16          6      OS                                       |
| _ 1      MUSCA17          7      OS                                       |
| _ 1      MUSCA18          8      OS                                       |
| _ 1      MUSCA19          9      OS                                       |
| / 2      MUSCA2A          A      CF    MUSCA2A test CF partition          |
| **************************** Bottom of data ***************************** |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 14-25 Processors: Define Access List, ICP

Because more partitions than the selected two are defined, the next panel `Define Candidate List` appears. For this example, we do not add any partitions in the access list. Therefore, press Enter.

HCD returns to the Channel Path List, which shows you the CHPID that was defined (see Figure 14-26).

```
                              Channel Path List       Row 1 of 5 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E8    124/1 CL5   SPAN  __  __ __   Y   No  FC 0433 Coupling Express LR
_ ED    144/2 CL5   SPAN  __  __ __   Y   No  FC 0433 Coupling Express LR
_ EE    _____ ICP   SPAN  __  __ __   N   No  ICP Internal Coupling Peer
```

Figure 14-26 Processors: Channel Path List, ICP

7. Define another ICP CHPID as EF to the same LPARs (see Figure 14-27).

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

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ E0    12/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E4    16/1  CS5   SPAN  __  __ __   Y   No  FC 0172 ICA SR 2 Links
_ E8    124/1 CL5   SPAN  __  __ __   Y   No  FC 0433 Coupling Express LR
_ ED    144/2 CL5   SPAN  __  __ __   Y   No  FC 0433 Coupling Express LR
_ EE    _____ ICP   SPAN  __  __ __   N   No  ICP Internal Coupling Peer
_ EF    _____ ICP   SPAN  __  __ __   N   No  ICP Internal Coupling Peer
```

Figure 14-27 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.

Creating a CF link using ICP CHPIDs includes the following considerations:

- Each of the logical ICP links can support only one CHPID at each end of the link.
- How many ICP CF links are required to provide enough coupling bandwidth.
- Which z/OS LPARs on the same processors need access to the CF LPARs.
- Our example connects two ICP CHPIDs (E4 and E5) on the same processor.

To define a Coupling Facility 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 includes 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 from which you want to connect, or in the Command Line enter `L EE`. In our example, we use `EE`.

5. Enter `f` next to the CHPID definition (see Figure 14-28), and press Enter.

```
                              Channel Path List       Row 5 of 6 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca
Configuration mode . : LPAR
Channel Subsystem ID : 2

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f EE    _____ ICP   SPAN  __  __ __    N   No  ICP Internal Coupling Peer
_ EF    _____ ICP   SPAN  __  __ __    N   No  ICP Internal Coupling Peer
```

Figure 14-28 Processors: CF Channel Path Connectivity List, ICP

6. HCD prompts you to specify the second ICP CHPID to which you want to connect. Scroll through the Channel Path List until you find the second ICP CHPID you want to connect to, or in the Command Line, enter `L EF`. In our example, we use `EF`.
7. Enter `p` next to the CHPID definition (see Figure 14-29) and press Enter.

```
                     CF Channel Path Connectivity List              Row 5 of
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ EE        Y  ICP  SPAN N
p EF        Y  ICP  SPAN N
```

Figure 14-29 Processors: CF Channel Path Connectivity List, ICP

8. HCD prompts you to specify to which CF channel path to connect. Complete the following updates (see Figure 14-30) and press Enter:

   – `Destination processor ID` to `MUSCA`
   – `Destination channel subsystem ID` to `2`
   – `Destination channel path ID` to `EE`

```
*---------------- Connect to CF Channel Path -----------------*
|                                                             |
|                                                             |
| Specify the following values.                               |
|                                                             |
| Source processor ID  . . . . . : MUSCA                      |
| Source channel subsystem ID  . : 2                          |
| Source channel path ID . . . . : EF                         |
| Source channel path type . . . : ICP                        |
|                                                             |
| Destination processor ID . . . . . . MUSCA___  +            |
| Destination channel subsystem ID . . 2   +                  |
| Destination channel path ID  . . . . EE  +                  |
|                                                             |
| Timing-only link . . . . . . . . . . No                     |
|                                                             |
*-------------------------------------------------------------*
```

Figure 14-30   Processors: Connect to CF Channel Path, ICP

9. HCD checks available Control Unit 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 ICP CHPID (see Figure 14-31) 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 . . . . . . . : MUSCA                                 |
| Channel subsystem ID . . . : 2                                     |
| Channel path ID  . . . . . : EF         Operation mode . . : SPAN  |
| Channel path type  . . . . : ICP                                   |
|                                                                    |
| Control unit number  . . . . FFFD  +                               |
|                                                                    |
| Device number  . . . . . . . FFBD                                  |
| Number of devices  . . . . . 7                                     |
|                                                                    |
*--------------------------------------------------------------------*
```

Figure 14-31   Processors: Add CF Control Unit and Devices, ICP

10. HCD provides suggested Control Unit numbers and Device addresses for the first ICP CHPID (see Figure 14-32). Observe that the Control unit 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 . . . . . . . : MUSCA                                  |
| Channel subsystem ID . . . : 2                                      |
| Channel path ID  . . . . . : EE         Operation mode . . : SPAN   |
| Channel path type  . . . . : ICP                                    |
|                                                                     |
| Control unit number  . . . . FFFD  +                                |
|                                                                     |
| Device number  . . . . . . . FFA5                                   |
| Number of devices  . . . . . 7                                      |
|                                                                     |
|                                                                     |
*---------------------------------------------------------------------*
```

Figure 14-32 Processors: Add CF Control Unit and Devices, ICP

HCD provides a summary of the Source and Destination CF Connectivity List (see Figure 14-33).

```
                     CF Channel Path Connectivity List              Row 5 of 6
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : MUSCA     Musca
Source channel subsystem ID . : 2
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ EE        Y  ICP  SPAN N     MUSCA.2    EF        Y  ICP  SPAN   CFP   7
_ EF        Y  ICP  SPAN N     MUSCA.2    EE        Y  ICP  SPAN   CFP   7
```

Figure 14-33 Processors: Add CF Channel Path Connectivity List, ICP

15

Adding PCIe devices

This chapter describes the steps to define ISM, RoCE-2, zEDC, and zHyperLink Peripheral Component Interconnect Express (PCIe) functions, features, and devices.

It includes a list of these potential configuration items and a short description on how to complete each of them by using Hardware Configuration Definition (HCD) and an overview of PCIe functions.

Note: Not all of the configuration items that are described in this chapter are necessarily required for your installation. Also, the examples that are presented are not intended to be exhaustive.

This chapter includes the following topics:

15.1 Defining more I/O by using HCD

When defining new I/O components in an IODF, specific definitions, such OSCONFIGs, partitions, FICON Switches, Control Units, and Devices, must be made up front. After these items are defined, the following connections can be made:

- Defining an ISM PCIe function
- Defining a RoCE-2 PCIe function
- Defining a zEDC EXPRESS PCIe function
- Defining a zHyperLink PCIe function

These I/O definitions use HCD to demonstrate the examples. The examples continue by using the example that was created in Chapter 5, "Production IODF and setting up the CPC" on page 85 (for example, `ITSO.IODF78.WORK`).

15.2 PCIe feature definitions

This section provides a brief overview of the `Function` statement and describes defining ISM, RoCE-2, zEDC, and zHyperLink PCIe features.

15.2.1 Overview

Starting with processor type 2827, PCIe adapters that are attached to a system can provide the operating system with various so-called PCIe functions to be used by entitled logical partitions (LPARs).

HCD supports the following components:

- Internal Shared Memory PCIe Adapter (ISM). A virtual PCIe adapter for which a virtual channel identifier (VCHID) must be defined.
- Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE). PCIe functions of type RoCE and RoCE-2 can be assigned to external physical networks by specifying corresponding PNET IDs.
- zEDC-Express. For PCIe functions of type zEDC-Express, a virtual function number 1 - 15 must be specified.
- Regional Crypto Enablement (RCE). This PCIe function type is used for vendor crypto adapters that are approved by IBM. This type not in the scope of this book.

Note: The support of virtual functions, the allowed range of virtual functions, and support of PNETIDs depends on the processor type and support level. For more information, see *Input/Output Configuration Program User's Guide*, SB10-7172. HCD offers prompts for virtual functions and ensures that the validation rules are fulfilled.

HCD provides dialogs to define, change, delete, and view PCIe functions, and to control which LPARs can access which PCIe functions.

PCIe adapters offer new functionality to systems running on processor 2827 and newer. HCD introduces a new dialog where users can define PCIe functions, assign them to LPARs, and activate them by using IOCP or dynamically.

HCD also provides the following new 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.

HCD supports the new I/O configuration statement FUNCTION for defining and configuring PCIe functions.

In the CPC, the input/output (I/O) subsystem, which controls channel operations, requires specific data about the hardware I/O configuration.

To define PCIe functions for the I/O subsystem, you must specify the following components:

- Logical partitions
- PCIe adapter functions on the CPC and their assignment to logical partitions.

A PCIe function is defined by a unique identifier, which is the function ID (FID). Each function specifies a function type and a channel identifier CHID. Multiple functions can be specified to the same CHID value that is provided that each of these functions defines a unique virtual function VF number when defining a PCIe function. Example 15-1 shows definitions for a zHyperlink card for Port 1 and multiple virtual functions, each assigned to a specific LPAR.

Example 15-1 zHyperlink definitions

```
***********************************************************************
*
      RESOURCE PART=((CSS(0),(Z01,1),(Z02,2),(Z03,3),(Z04,4),(CF01,8), *
               (CF02,9)),                                              *
               (CSS(1),(Z11,1),(Z12,2),(CF11,8),(CF12,9)),             *
               (CSS(2),(Z21,1),(Z22,2),(CF21,8)),                      *
               (CSS(3),(Z31,1),(Z32,2),(CF31,8)),                      *
               (CSS(4),(Z41,1),(Z42,2),(LX41,3),(LX42,4),(CF41,8)),    *
               (CSS(5),(Z51,1),(Z52,2),(CF51,8)))
*
***********************************************************************
****  SYNCH_IO
***************************************************************
**
 FUNCTION  FID=1200,PCHID=124,TYPE=HYL,PORT=1,VF=1,                    *
               PART=((Z03),(=))
 FUNCTION  FID=1201,PCHID=124,TYPE=HYL,PORT=1,VF=2,                    *
               PART=((Z04),(=))
 FUNCTION  FID=1202,PCHID=124,TYPE=HYL,PORT=1,VF=3,                    *
               PART=((Z11),(=))
 FUNCTION  FID=1203,PCHID=124,TYPE=HYL,PORT=1,VF=4,                    *
               PART=((Z21),(=))
 FUNCTION  FID=1204,PCHID=124,TYPE=HYL,PORT=1,VF=5,                    *
               PART=((Z31),(=))
 FUNCTION  FID=1205,PCHID=124,TYPE=HYL,PORT=1,VF=6,                    *
               PART=((Z41),(=))
 FUNCTION  FID=1206,PCHID=124,TYPE=HYL,PORT=1,VF=7,                    *
               PART=((Z42),(=))
 FUNCTION  FID=1207,PCHID=124,TYPE=HYL,PORT=1,VF=8,                    *
               PART=((Z51),(=))
```

The applicable functions to the various function types are listed Table 15-1.

Table 15-1 Keyword applicability for functions

FUNCTION	Function type					
	HYL	ISM	RCE	ROCE	ROC2	ZEDC
PCHID	Yes	No	Yes	Yes	Yes	Yes
VCHID	No	Yes	No	No	No	No
PNETID	No	Yes	No	Yes	Yes	No
PORT	Yes	No	No	No	Yes	No
VF	Yes	Yes	No	Yes	Yes	Yes

For more information about the maximum values for each machine type, see *IBM Z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7172.

15.2.2 Defining an ISM PCIe function

The z14 ZR1 supports Internal Shared Memory (ISM) virtual PCIe (vPCIe) device to enable optimized cross-LPAR TCP communications by using a socket-based direct memory access (DMA), the Shared Memory Communications - Direct Memory Access (SMC-D).

SMC-D uses a virtual PCIe adapter and is configured as is a physical PCIe device. Up to 32 ISM adapters are available, each with a unique Physical Network ID per CPC.

In this example, we define the following items:

- CHID=7F1 to Function IDs 0040 (VF=1) and 0041 (VF=2) on CPC = MUSCA
- CHID=7F2 to Function IDs 0050 (VF=1) and 0051 (VF=2) on CPC = MUSCA

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 (CETUS) to which you want to define the ISM functions and press Enter (see Figure 15-1).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
f MUSCA    3907     ZR1      LPAR  007A883907 Musca
```

Figure 15-1 Processor List: Adding PCIe functions to a processor, ISM

3. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (see Figure 15-2).

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
```

Figure 15-2 PCIe Function List: Adding PCIe functions to a processor, ISM

4. Complete the following updates (see Figure 15-3) and press Enter:
   - `Function ID` to `0040`.
   - `Type` to `ISM`.
   - `Channel ID` to `7F1`.
   - `Virtual Function ID` to `1`.
   - `Description` to the description that you want.

```
*---------------------------- Add PCIe Function ----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID  . . . . : MUSCA        Musca                                |
|                                                                           |
| Function ID . . . . . . 0040                                              |
| Type  . . . . . . . . . ISM           +                                   |
|                                                                           |
| Channel ID . . . . . . . . . . . 7F1  +                                   |
| Port . . . . . . . . . . . . . . _    +                                   |
| Virtual Function ID  . . . . . . 1__  +                                   |
| Number of virtual functions  . . 1                                        |
| UID  . . . . . . . . . . . . . . ____                                     |
|                                                                           |
| Description  . . . . . . . . . . ISM FID=0040 VFID=1 CHID=7F1____         |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 15-3 PCIe Function List: Adding PCIe functions to a processor, ISM

5. Update the `Physical network ID` to `PERFNET` value in the Add/Modify Physical Network IDs panel (see Figure 15-4) 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  . . ________________                               |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 15-4 Add/Modify Physical Network IDs: Adding network ID, ISM

6. Select the required Access LPAR for Function access list. In our example, we use LPAR `MUSCA11 (OS)`. Press Enter (see Figure 15-5).

```
*--------------------------- Define Access List ----------------------------*
|                                                              Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                            |
| Select one partition for the access list.                                 |
|                                                                            |
| Function ID  . . . . : 0040                                               |
|                                                                            |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA11          1      OS                                       |
| _ 1      MUSCA12          2      OS                                       |
| _ 1      MUSCA13          3      OS                                       |
| _ 1      MUSCA14          4      OS                                       |
| _ 1      MUSCA15          5      OS                                       |
| _ 1      MUSCA16          6      OS                                       |
| _ 1      MUSCA17          7      OS                                       |
| _ 1      MUSCA18          8      OS                                       |
| _ 1      MUSCA19          9      OS                                       |
| _ 2      MUSCA2A          A      CF    MUSCA2A test CF partition          |
| **************************** Bottom of data ***************************** |
|                                                                            |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 15-5 Define Access List: Selecting partition for Function access, ISM

7. Select any Candidate LPARs for Function access list. (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 15-6).

```
                              PCIe Function List     Row 1 of 21 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 0040   7F1    _   1    ISM           ____  ISM FID=0040 VFID=1 CHID=7F1
```

Figure 15-6 PCIe Function List: Function now created, ISM

8. Define the other Function IDs according to the example thus far (see Figure 15-7).

```
                              PCIe Function List     Row 1 of 24 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  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
_ 0050   7F2    _   1    ISM           ____  ISM FID=0050 VFID=1 CHID=7F2
_ 0051   7F2    _   2    ISM           ____  ISM FID=0051 VFID=2 CHID=7F2
```

Figure 15-7 PCIe Function List: All Functions now created, ISM

15.2.3 Defining a RoCE-2 PCIe function

RDMA over Converged Ethernet (RoCE) uses a new PCIe hardware card for the z14 and is named 10GbE RoCE Express (FC 0412).

To use RoCE hardware functions, the Function IDs, Virtual Function IDs, and Physical Network IDs must be defined in HCD to the PCHID that is assigned to the RoCE hardware cards that are installed in the processor.

Similar to defining a channel-path identifier (CHPID) to a physical channel ID (PCHID) for FICON and OSA type channels, Function IDs and Virtual Function IDs are assigned RoCE PCHIDs.

For more information about the PCHID and Resource Group (RG), see the PCHID report for the processor (see Example 15-2).

Example 15-2 PCHID Report: RoCE information from PCHID report

```
 Machine: 3907-ZR1  NEW1
  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
  Source          Drwr  Slot  F/C    PCHID/Ports or AID             Comment
  A09/LG02/J01    A14B  02    0412   100/D1D2            RG1
  A09/LG10/J01    A01B  12    0412   160/D1D2            RG2
  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
Legend:
  Source  Book Slot/Fanout Slot/Jack
  RG1     Resource Group 1
  0412    10GbE  RoCE Express
  RG2     Resource Group 2
```

In this example, we define the following items:

- PCHID=100 to Function IDs 00A2 (VF=3) Port=1and 00A3 (VF=4) Port=2 on CPC = MUSCA to Physical Network ID 1 = PERFNET
- PCHID=1BC to Function IDs 0010 (VF=1) Port=1and 0011 (VF=2) Port=2 on CPC = MUSCA to Physical Network ID 1 = PERFNET

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 (`CETUS`) to which you want to define the RoCE-2 functions and press Enter (see Figure 15-8).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
f MUSCA    3907     ZR1      LPAR  007A883907 Musca
```

Figure 15-8 Processor List: Adding PCIe functions to a processor, RoCE-2

3. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (see Figure 15-9).

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA       Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
```

Figure 15-9 PCIe Function List: adding PCIe functions to a processor, RoCE-2

4. Complete the following updates (see Figure 15-10) and press Enter:
   – `Function ID` to 00A2.
   – `Type` to ROCE-2.
   – `Channel ID` to 100.
   – `Port` to 1.
   – `Virtual Function ID` to 3.
   – `Description` to the description that you want.

```
*---------------------------- Add PCIe Function ----------------------------*
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID  . . . . : MUSCA       Musca                                  |
|                                                                            |
| Function ID . . . . . . 00A2                                               |
| Type  . . . . . . . . . ROCE-2       +                                     |
|                                                                            |
| Channel ID . . . . . . . . . . . 100  +                                    |
| Port . . . . . . . . . . . . . . 1    +                                    |
| Virtual Function ID  . . . . . . 3__  +                                    |
| Number of virtual functions  . . 1                                         |
| UID  . . . . . . . . . . . . . . ____                                      |
|                                                                            |
| Description  . . . . . . . . . . RoCE-2 FID=00A2 VFID=3 CHID=100_         |
|                                                                            |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 15-10 PCIe Function List: Adding PCIe functions to a processor, RoCE-2

5. Update the `Physical network ID` to `PERFNET` value in the Add/Modify Physical Network IDs panel (see Figure 15-11) 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  . . ________________                              |
|                                                                         |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 15-11 Add/Modify Physical Network IDs: Adding network ID, RoCE-2

6. Select the required Access LPAR for Function access list (see Figure 15-12). In our example, we use LPAR CETUS21 (OS). Press Enter.

```
*-------------------------- Define Access List ---------------------------*
|                                                          Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                         |
| Select one partition for the access list.                               |
|                                                                         |
| Function ID  . . . . : 00A2                                             |
|                                                                         |
| / CSS ID Partition Name   Number Usage Description                      |
| / 1      MUSCA11          1      OS                                     |
| _ 1      MUSCA12          2      OS                                     |
| _ 1      MUSCA13          3      OS                                     |
| _ 1      MUSCA14          4      OS                                     |
| _ 1      MUSCA15          5      OS                                     |
| _ 1      MUSCA16          6      OS                                     |
| _ 1      MUSCA17          7      OS                                     |
| _ 1      MUSCA18          8      OS                                     |
| _ 1      MUSCA19          9      OS                                     |
| _ 2      MUSCA2A          A      CF    MUSCA2A test CF partition        |
| **************************** Bottom of data ***************************** |
|                                                                         |
|                                                                         |
*-------------------------------------------------------------------------*
```

Figure 15-12 Define Access List: Selecting partition for Function access, RoCE-2

7. Select any Candidate LPARs for Function access list. In our example, we do *not* select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel in which you can see the Function now defined (see Figure 15-13).

```
                              PCIe Function List     Row 7 of 24 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 00A2   100    1   3    ROCE-2        ____  RoCE-2 FID=00A2 VFID=3 CHID=100
```

Figure 15-13 PCIe Function List: Function now created, RoCE-2

8. Define the other Function IDs according to the example thus far (see Figure 15-14).

```
                              PCIe Function List     Row 7 of 24 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 00A2   100    1   3    ROCE-2        ____  RoCE-2 FID=00A2 VFID=3 CHID=100
_ 00A3   100    2   4    ROCE-2        ____  RoCE-2 FID=00A3 VFID=4 CHID=100
_ 00B2   160    1   3    ROCE-2        ____  RoCE-2 FID=00B2 VFID=3 CHID=160
_ 00B3   160    2   4    ROCE-2        ____  RoCE-2 FID=00B3 VFID=4 CHID=160
```

Figure 15-14 PCIe Function List: All Functions now created, RoCE-2

15.2.4 Defining a zEDC EXPRESS PCIe function

The zEnterprise Data Compression (zEDC) uses a PCIe hardware card for the z14 ZR1called zEDC Express FC#0420.

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that was assigned to the zEDC hardware cards that are installed in the processor. Similar to defining a CHPID to a PCHID for FICON and OSA type channels, Function IDs and Virtual Function IDs are assigned zEDC PCHIDs.

This example shows only the definition process. For more information about how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

For more information about the PCHID and RG, see the PCHID report for the processor (see Example 15-3).

Example 15-3 PCHID Report: zEDC information from PCHID report

```
Machine: 3907-ZR1  NEW1
  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
  Source          Drwr  Slot  F/C    PCHID/Ports or AID            Comment
  A09/LG08/J01    A14B  12    0420   120                   RG2
  A09/LG04/J01    A01B  02    0420   140                   RG1
Legend:
  Source  Book Slot/Fanout Slot/Jack
  RG1     Resource Group 1
  RG2     Resource Group 2
  0420    zEDC Express
```

In this example, we define the following items:

- CHID=120 to Function IDs 00C1 (VF=2) on CPC = MUSCA
- CHID=140 to Function IDs 00D1 (VF=2) on CPC = MUSCA

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 (`CETUS`) to which you want to define the zEDC functions and press Enter (see Figure 15-15).

```
                              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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
f MUSCA    3907     ZR1      LPAR  007A883907 Musca
```

Figure 15-15 Processor List: Adding PCIe functions to a processor, zEDC

3. Enter `add` on the command line in the PCIe Function List panel (see Figure 15-16) to add a PCIe function.

```
                            PCIe Function List
Command ===> add________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA       Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
```

Figure 15-16 PCIe Function List: Adding PCIe functions to a processor, zEDC

4. Complete the following updates (see Figure 15-17) and press Enter:

   – `Function ID` to **00C1**.
   – `Type` to **ZEDC-EXPRESS**.
   – `Channel ID` to **120**.
   – `Virtual Function ID` to **1**.
   – `Description` to the description that you want.

```
*---------------------------- Add PCIe Function ----------------------------*
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID  . . . . : MUSCA        Musca                                |
|                                                                           |
| Function ID . . . . . . 00C1                                              |
| Type  . . . . . . . . . ZEDC-EXPRESS  +                                   |
|                                                                           |
| Channel ID . . . . . . . . . . . 120  +                                   |
| Port . . . . . . . . . . . . . . _    +                                   |
| Virtual Function ID  . . . . . . 2    +                                   |
| Number of virtual functions  . . 1                                        |
| UID  . . . . . . . . . . . . . . ____                                     |
|                                                                           |
| Description  . . . . . . . . . . zEDC FID=00C1 VFID=2 CHID=120            |
|                                                                           |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 15-17 PCIe Function List: Adding PCIe functions to a processor, zEDC

5. Select the required Access LPAR for Function access list. In our example, we use LPAR `MUSCA11 (OS)`. Then, press Enter (see Figure 15-18).

```
*--------------------------- Define Access List ----------------------------*
|                                                              Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one partition for the access list.                                 |
|                                                                           |
| Function ID  . . . . : 00C1                                               |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| / 1      MUSCA11          1      OS                                       |
| _ 1      MUSCA12          2      OS                                       |
| _ 1      MUSCA13          3      OS                                       |
| _ 1      MUSCA14          4      OS                                       |
| _ 1      MUSCA15          5      OS                                       |
| _ 1      MUSCA16          6      OS                                       |
| _ 1      MUSCA17          7      OS                                       |
| _ 1      MUSCA18          8      OS                                       |
| _ 1      MUSCA19          9      OS                                       |
| _ 2      MUSCA2A          A      CF    MUSCA2A test CF partition          |
| **************************** Bottom of data ***************************** |
|                                                                           |
|                                                                           |
*---------------------------------------------------------------------------*
```

Figure 15-18 Define Access List: selecting partition for Function access, zEDC

6. Select any Candidate LPARs for Function access list. In our example, we do not select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel (see Figure 15-19) where you can see that the Function is now defined.

```
                              PCIe Function List    Row 14 of 23 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 00C1   120    _   2    ZEDC-EXPRESS  ____  zEDC FID=00C1 VFID=2 CHID=120
```

Figure 15-19 PCIe Function List: Function now created, zEDC

7. Define the other Function IDs according to this example (see Figure 15-20).

```
                              PCIe Function List    Row 14 of 24 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 00C1   120    _   2    ZEDC-EXPRESS  ____  zEDC FID=00C1 VFID=2 CHID=120
_ 00D1   140    _   2    ZEDC-EXPRESS  ____  zEDC FID=00D1 VFID=2 CHID=140
```

Figure 15-20 PCIe Function List: All Functions now created, zEDC

15.2.5 Defining a zHyperLink PCIe function

The zHyperLink Express is a direct connect short distance IBM Z I/O feature that is designed to work with a High-Performance FICON SAN infrastructure.

IBM zHyperLink dramatically reduces latency by interconnecting the z14 ZR1 CPC directly to the I/O bay of the DS8880.

zHyperLink uses a PCIe hardware card for the z14 and the z14 ZR1 called zHyperLink Express (FC 0431).

Two ports are available per feature and a maximum of 16 features can be installed on a z14 ZR1. Up to 127 Virtual Functions can be defined per CHID.

In this example, we define the following items:

- CHID=108 to FID = 0102, VFIDs = 3, Port = 1, on CPC = MUSCA
- CHID=108 to FID = 0103, VFIDs = 4, Port = 2, on CPC = MUSCA
- CHID=168 to FID = 0202, VFIDs = 3, Port = 1, on CPC = MUSCA
- CHID=168 to FID = 0203, VFIDs = 4, Port = 2, on CPC = MUSCA

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 (`CETUS`) to which you want to define the zHyperLink functions and press Enter (see Figure 15-21).

```
                                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
_ CETUS    3906     M03      LPAR  0EE0F73906 Cetus
_ LEPUS    2965     N20      LPAR  0BB4B72965 Lepus
f MUSCA    3907     ZR1      LPAR  007A883907 Musca
```

Figure 15-21 Processor List: Adding PCIe functions to a processor, zHyperLink

2. To add a PCIe function, enter `add` on the command line in the PCIe Function List panel (see Figure 15-22).

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
```

Figure 15-22 PCIe Function List: Adding PCIe functions to a processor, zHyperLink

3. Complete the following updates (see Figure 15-23) and press Enter:
   - `Function ID` to 102.
   - `Type` to ZHYPERLINK.
   - `Channel ID` to 108.
   - `Port ID` to 1.
   - `Virtual Function ID` to 3.
   - `Description` to the description that you want.

```
*---------------------------- Add PCIe Function ----------------------------*
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID  . . . . : MUSCA        Musca                                 |
|                                                                            |
| Function ID . . . . . . 0102                                               |
| Type  . . . . . . . . . ZHYPERLINK    +                                    |
|                                                                            |
| Channel ID . . . . . . . . . . . 108  +                                    |
| Port . . . . . . . . . . . . . . 1    +                                    |
| Virtual Function ID  . . . . . . 3__  +                                    |
| Number of virtual functions  . . 1                                         |
| UID  . . . . . . . . . . . . . . ____                                      |
|                                                                            |
| Description  . . . . . . . . . . zHYP FID=0102 VFID=3 CHID=108___          |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 15-23 PCIe Function List: Adding PCIe functions to a processor, zHyperLink

4. Select the required Access LPAR for Function access list. In our example, we use LPAR `CETUS21 (OS)`. Press Enter (see Figure 15-24).

```
*--------------------------- Define Access List ----------------------------*
|                                                               Row 22 of 31 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                            |
| Select one partition for the access list.                                  |
|                                                                            |
| Function ID  . . . . : 0102                                                |
|                                                                            |
| / CSS ID Partition Name   Number Usage Description                         |
| / 1      MUSCA11          1      OS                                        |
| _ 1      MUSCA12          2      OS                                        |
| _ 1      MUSCA13          3      OS                                        |
| _ 1      MUSCA14          4      OS                                        |
| _ 1      MUSCA15          5      OS                                        |
| _ 1      MUSCA16          6      OS                                        |
| _ 1      MUSCA17          7      OS                                        |
| _ 1      MUSCA18          8      OS                                        |
| _ 1      MUSCA19          9      OS                                        |
| _ 2      MUSCA2A          A      CF    MUSCA2A test CF partition          |
| *************************** Bottom of data **************************** |
|                                                                            |
|                                                                            |
*----------------------------------------------------------------------------*
```

Figure 15-24 Define Access List: Selecting partition for Function access, zHyperLink

5. Select any Candidate LPARs for Function access list. In our example, we do *not* select any candidate LPARs. Press Enter.

   HCD returns to the PCIe Function List panel in which you can see the Function now defined (see Figure 15-25).

```
                              PCIe Function List    Row 19 of 21 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 0102   108    1   3    ZHYPERLINK    ____  zHYP FID=0102 VFID=3 CHID=108
```

Figure 15-25 PCIe Function List: Function now created, zHyperLink

6. Define the other Function IDs according to the example thus far (see Figure 15-26).

```
                              PCIe Function List    Row 13 of 24 More:
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : MUSCA        Musca



/ FID    CHID+  P+  VF+  Type+         UID   Description
_ 0102   108    1   3    ZHYPERLINK    ____  zHYP FID=0102 VFID=3 CHID=108
_ 0103   108    2   4    ZHYPERLINK    ____  zHYP FID=0103 VFID=4 CHID=108
_ 0202   168    1   3    ZHYPERLINK    ____  zHYP FID=0202 VFID=3 CHID=168
_ 0203   168    2   4    ZHYPERLINK    ____  zHYP FID=0203 VFID=4 CHID=168
```

Figure 15-26 PCIe Function List: All Functions now created, zHyperLink

16

Configuring a z14 ZR1 server by using Dynamic Partition Manager

This chapter describes how to configure a z14 ZR1 CPC by using Dynamic Partition Manager (DPM). DPM is a management interface that simplifies the configuration and administration of an IBM Z CPC.

This chapter includes the following topics:

16.1 Introduction

DPM is a resource management and operational environment that provides a simplified approach to configuring and managing IBM Z servers[1]. DPM reduces the barriers to the adoption of IBM Z for new and existing customers.

The DPM implementation provides built-in integrated capabilities that allow advanced virtualization management on IBM Z servers. With DPM, customers can use their Linux and virtualization skills while getting the full value of IBM Z hardware's robustness and rich security features in a workload optimized environment.

DPM provides facilities to define and run virtualized computing systems by using a firmware-managed environment that coordinates the physical system resources that are shared by the partitions[2]. The partitions' resources include processors, memory, network, storage, Crypto, and Accelerators.

DPM is not an extra hypervisor for IBM Z servers. DPM uses the PR/SM hypervisor infrastructure and provides an intelligent and unified interface that allows customers to define, use, and operate the platform virtualization with little or no IBM Z experience.

> **Note:** When IBM z14 ZR1 servers are set to run in DPM mode, the following components are supported:
>
> - Linux virtual servers running in a partition
> - The KVM hypervisor[a] for Linux guests
> - z/VM with Linux guests
> - Virtual appliances that use the IBM Secure Service Container (SSC) framework
>
> a. Available with Linux distributions.

The DPM infrastructure is shown in Figure 16-1.

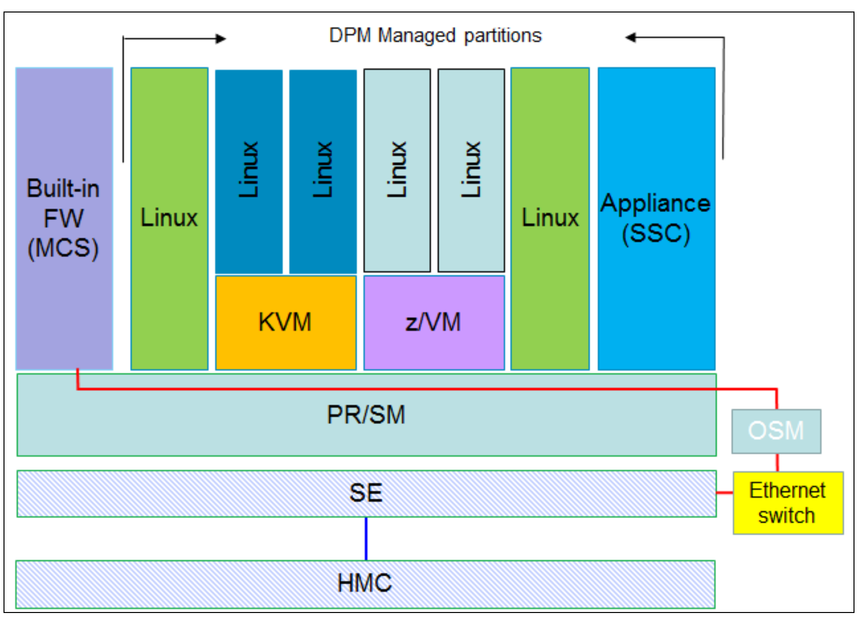

Figure 16-1 High-level overview of DPM implementation

[1] IBM LinuxONE servers also can be managed by using DPM.

[2] DPM uses the term *partition*, which is the same as logical partition (LPAR).

The firmware partition (similar to the PCIe support partition, which is also known as the master control services [MCS] partition), along with the Support Element (SE), provides instrumentation to create and manage the Linux native partitions (used also for deploying kernel-based virtual machine (KVM) code), partitions that run the IBM z/VM hypervisor, or partitions that run appliances that use the Secure Service Container framework.

The connectivity from the SE to the MCS is provided through the internal management network by two OSA-Express 1000BASE-T that are acting as management adapters (OSM).

16.1.1 Prerequisites

This section provides an overview for DPM prerequisites and considerations for supported features. For more information, see the *IBM Dynamic Partition Manager (DPM) Guide,* SB10-7170.

Prerequisites

The DPM is a mode of operation that requires IBM z14 M0x (Machine Type 3906), IBM z14 ZR1 (Machine Type 3907), IBM z13 (Machine Type 2964 - Driver Level 27), or z13s (Machine Type 2965) CPCs.

> **Feature enablement:** DPM is a feature code (FC) 0016 that can be selected during the machine order process. FC 0016 requires one pair of OSA-Express 1000BASE-T adapters, which must be included in the configuration.
>
> When the Z CPC is managed by using DPM, the OSA-Express 1000BASE-T features are dedicated for DPM configuration and *cannot* be used for any other purpose.

16.1.2 Planning considerations

In this section, we provide considerations for installing a new system (create a configuration from scratch) and for switching a system from “traditional” management to use DPM.

We also describe I/O features and other configuration aspects, such as time synchronization and resource allocation.

New system

If the system you are configuring is new (no pre-existing configuration to be saved or migrated), start configuring the system as described in 16.1.4, “Enabling DPM” on page 391.

> **Important:** Consider the following points:
>
> - A CPC in DPM mode cannot be part of an Ensemble that is managed by Unified Resource Manager. The HMC that is used to enable the CPC in DPM mode must *not* be an Ensemble HMC (Primary or Backup Ensemble HMC).
> - All definitions that are made for the CPC (if any) before the DPM mode is activated are saved and can be brought back if you choose to revert to standard PR/SM mode. However, when switching the CPC into standard PR/SM mode, any definitions that are made with the CPC in DPM mode are discarded (not saved).

Previously configured system

If you are converting a system (configured and managed by using standard or traditional tools) to the use of DPM, you must check the configuration of the system and verify whether you must save any configuration data. This process helps you to decide whether you should revert to "traditional" management tools.

16.1.3 Considerations for I/O features

The following I/O and special purpose features are supported for configuration by using DPM:

- OSA-Express features (all supported features)
- FICON Express (all supported features):
  - Fibre Channel Protocol (SCSI)
  - FICON (ECKD) DASD
- RoCE Express and RoCE Express2 features (as standard NIC, no SMC-R support)
- IBM zEDC Express (used as an accelerator)
- Crypto Express features (all supported features)

> **Important:** Not all channel types are supported for these features. For more information about supported channel types, see the DPM documentation.

The following features are *not* supported when DPM is used to configure and manage your system:

- Coupling features[3]:
  - ICA SR
  - Coupling Express Long Reach
- zHyperLink Express

> **Important:** Always consult with operating system support for the features' functionality availability.

Internal features

From the available internal capabilities of the z14 ZR1, HiperSockets are supported and configurable by using DPM.[4]

Time synchronization

If your environment requires time synchronization, the z14 ZR1 includes Server Time Protocol (STP) support. STP support is provided through FC 1021.

STP allows the Z CPC to synchronize its time to an external time source (ETS). At Power On Reset, the CPC retrieves the Time Of Day (TOD) from the SE. Further on, the time of the CPC can be maintained in sync with an ETS by using the SE Network Time Protocol (NTP) client (with STP configured) and STP feature.

---

[3] InfiniBand coupling features are not supported on z14 ZR1. Although these features might be supported on systems that support DPM, the InfiniBand features are *not* configurable or usable when the CPC is managed by DPM.

[4] Other internal features include Internal Shared Memory (ISM) communications that are used for SMC-D, Internal Coupling (IC), and Virtual Flash Memory (VFM). These features are not supported in DPM mode.

When configured, STP creates a Coordinated Timing Network (CTN) that is used for maintaining CPCs in sync with an ETS.

When in DPM mode, a CPC can be configured as a *Single Server CTN*. Through STP, time synchronization information is provided to the operating systems that is running in the partitions (LPARs) on the CPC.

Operating system support is required for the use of the STP feature. Linux on Z and z/VM support time synchronization by using STP.

STP tip: STP provides timing information to all partitions in a CPC. If the CPC uses STP, all active partitions can maintain time synchronization to an external source without requiring external timing source access (through a network connection).

Moreover, the IBM Z CPC is equipped with Pulse Per Second input to increase time synchronization accuracy from 100 ms (standard NTP accuracy) to 10 microseconds when a time source (NTP server) is used that also includes a Pulse Per Second output (PPS).

For more information, see the following publications:

- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Implementation Guide*, SG24-7281
- *Server Time Protocol Recovery Guide*, SG24-7380

Processor considerations

The following processor characterization types can be configured and used for *customer workload* on a CPC that is managed by using DPM:

- Central Processor (CP)
- Integrated Facility for Linux (IFL)

Important: Simultaneous Multi-Threading (SMT) is not supported on CPs. As of this writing, only IFLs support SMT.

The following processor types *cannot* be configured when the CPC is managed by using DPM:

- IBM Z Integrated Information Processor (zIIP)
- Integrated Coupling Facility (ICF)

16.1.4 Enabling DPM

Enabling DPM is a disruptive action. The selection of DPM mode of operation is done by using a SE task that is named Enable Dynamic Partition Manager under the CPC Configuration menu, as shown in Figure 16-2 on page 392.

For more information about DPM enabling, see *IBM 3907 Installation Manual,* GC28-6973.

Note: During the machine installation[a] or in preparation for DPM (later), the IBM SSR connects the two OSA Express 1000BASE-T cables to the Ethernet Top of Rack (management) switches in the designated ports.

a. The DPM mode of operation setting is normally performed at machine installation time by the service support representative (SSR).

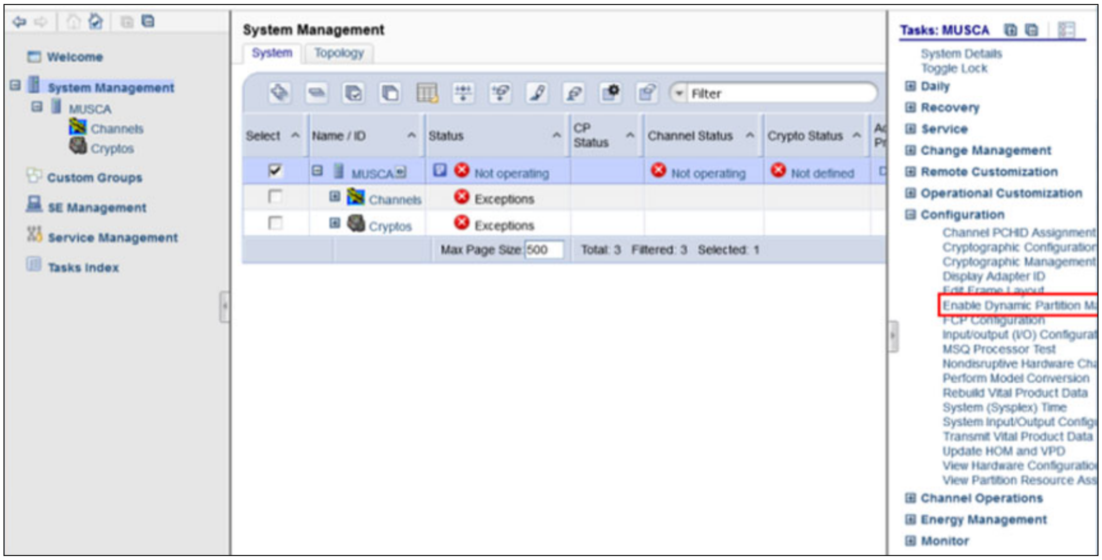

Figure 16-2 Enabling DPM mode of operation from the SE CPC configuration tasks

After the option is selected, a new window opens (see Figure 16-3) in which you must enter the two OSA Express 1000BASE-T ports that are designated for DPM management (connected to the two switches in the frame during the Z server installation).

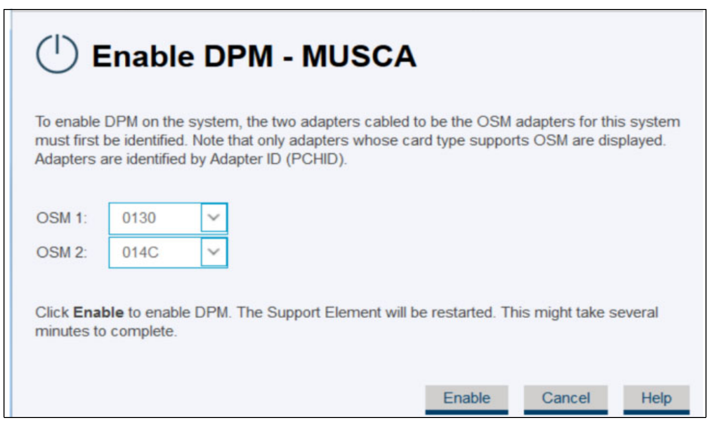

Figure 16-3 Entering the OSA ports that are used by the management network

After entering the OSA adapter port numbers that were cabled to the switches, click **Enable**.

The SE then restarts. When finished, the DPM mode becomes active and operational.

The DPM mode welcome window is shown in Figure 16-4. The three options at the bottom (*Getting Started*, *Guides*, and *Learn More*) include mouse-over functions that briefly describe their meaning or provide more functions.

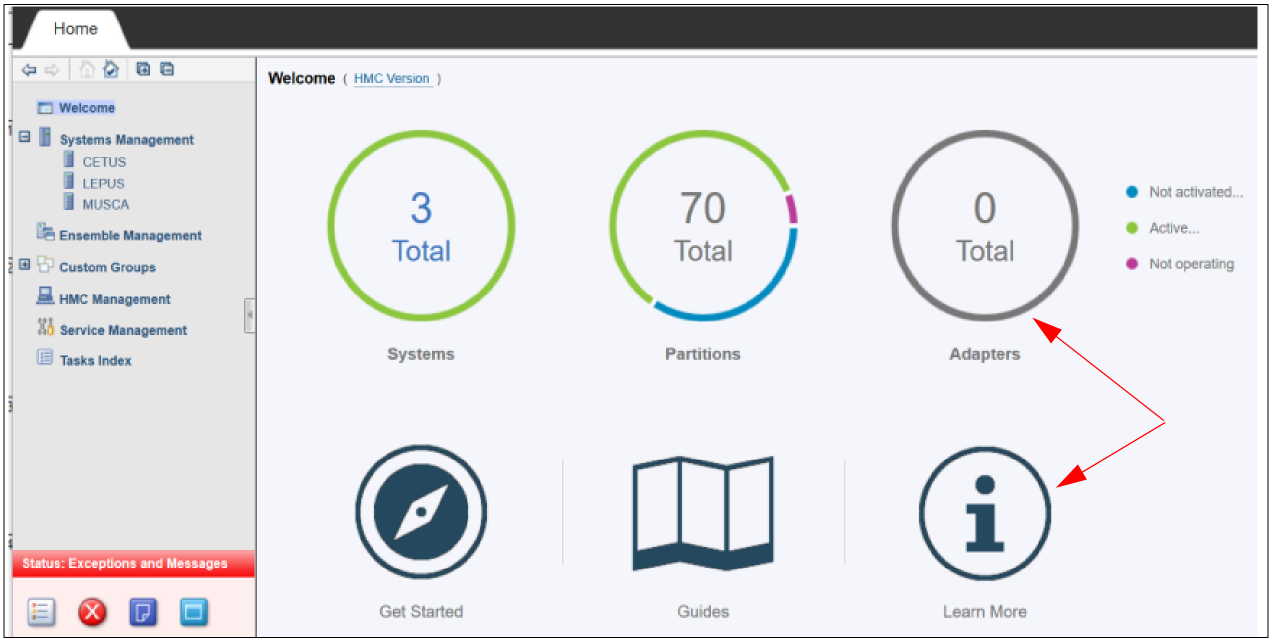

Figure 16-4 DPM mode welcome window

The welcome window that is shown in Figure 16-4 opens only when at least one HMC-defined CPC is active in DPM mode. This DPM mode can be identified by the presence of the Adapters information. Otherwise, the traditional HMC window is presented when you log on to the HMC (the items that are indicated by the arrows in Figure 16-4 are not displayed).

16.2 Configuring Server Time Protocol

Note: This configuration task is optional. You can skip this task if you do not plan to use STP for CPC time synchronization.

However, if you plan to use STP, this task is performed *before* partitions are configured and activated; otherwise, the partitions might need to be reactivated to acquire any timing configuration from STP if configured later.

STP configuration for a single server CTN configuration requires SE connectivity to an external NTP server (external time source). Alternatively, the HMC can be configured as an NTP server.

STP can be used if time synchronization for our environment is required. Although STP is an optional feature (FC 1021), it is highly recommended (for more information, see “Time synchronization” on page 390).

This section describes the procedure that we used in our environment to enable STP for a DPM-managed CPC.

16.2.1 Configuring External Time Source on the Support Element

Note: Ensure that you configured at least one NTP server to be accessible to the SE. Communication with the designated NTP server also can be encrypted. For more information about configuring encryption for the NTP server, see the HMC Version 2.14.0 documentation.

The STP configuration information can be checked by using the CPCs Support Element interface. In our example, we used the following process to check the configuration:

1. From the HMC interface, with a user that includes System Programmer (Sysprog) authority, we selected our CPC. Then, clicked **Recovery** → **Single Object Operations** (SOO). In the SOO interface, we select the CPC and then, click the **System (Sysplex) Time** task (see Figure 16-5).

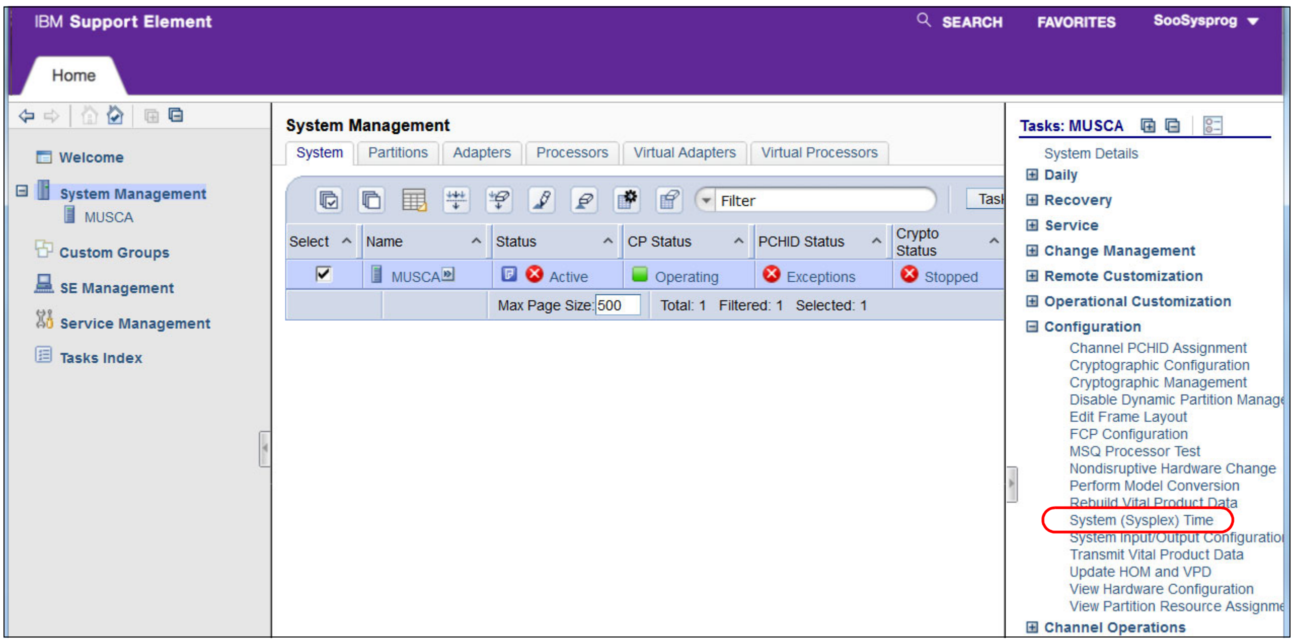

Figure 16-5 Accessing System (Sysplex) Time menus form the SE

The window that shown in Figure 16-6 opens. If the CPC is newly installed or STP was not yet configured, no Timing Network (CTN) information is available.

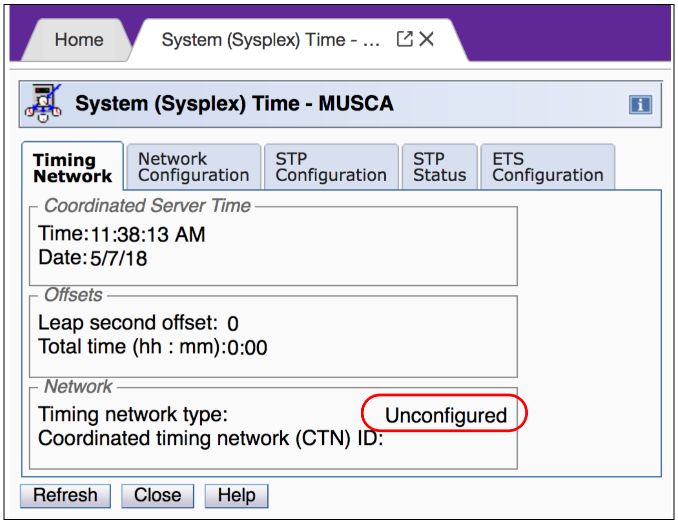

Figure 16-6 System (Sysplex) Time configuration information (SE view)

2. The ETS information must be configured. We click the **ETS Configuration** tab, as shown in Figure 16-7.

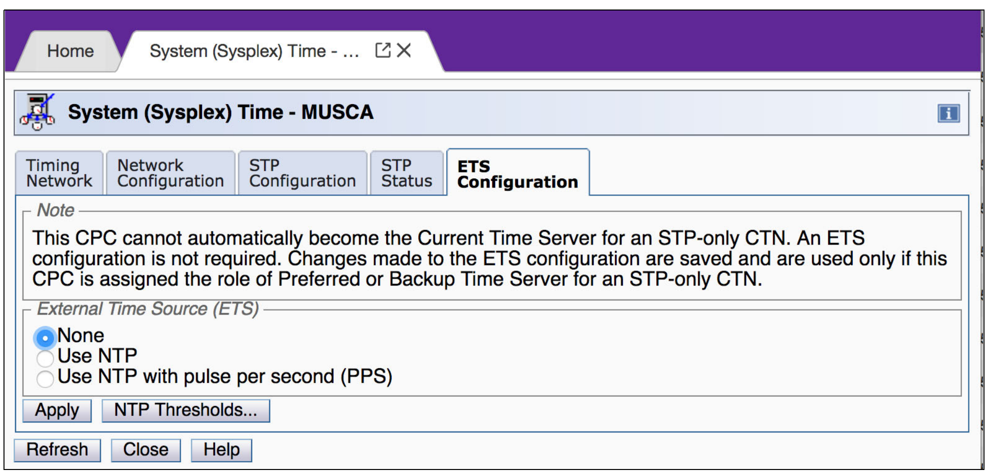

Figure 16-7 External Time Source tab (not configured)

3. We select **Use NTP** and enter the necessary information. Then, we check the NTP server connectivity (by clicking **Query**), as shown in Figure 16-8 on page 396. After the NTP servers are contacted, the NTP stratum and time source information are shown.

   The Adjust NTP threshold option can be used if your NTP servers might drift away from a lower NTP stratum level (through NTP reconfiguration).

> **Note:** The lower the NTP stratum value, the better the time synchronization accuracy for your system.

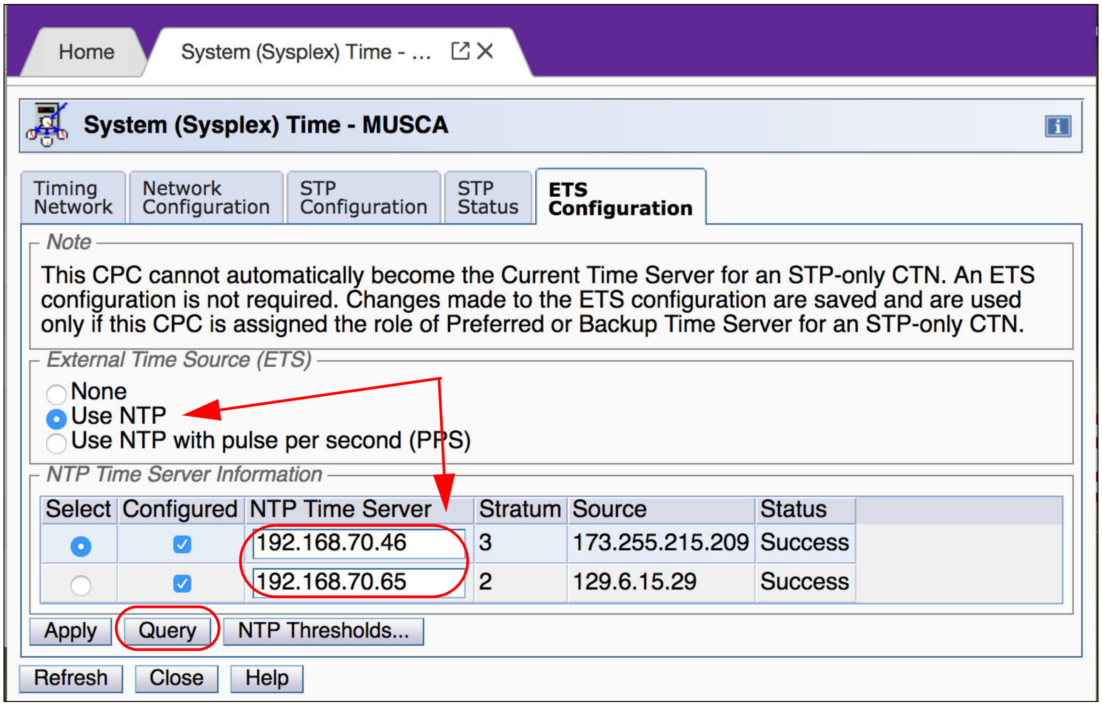

Figure 16-8 Entering ETS information

4. We click **Apply** to save ETS information. This information is used later to configure the CTN.
5. We close the System (Sysplex) Timer task and log off from the support element.

16.2.2 Configuring the Coordinated Timing Network (HMC)

We completed the following steps to configure the CTN for the CPC in DPM mode:

1. We continue STP configuration on the HMC by selecting the CPC and then clicking **Manage System Time**, as shown in Figure 16-9.

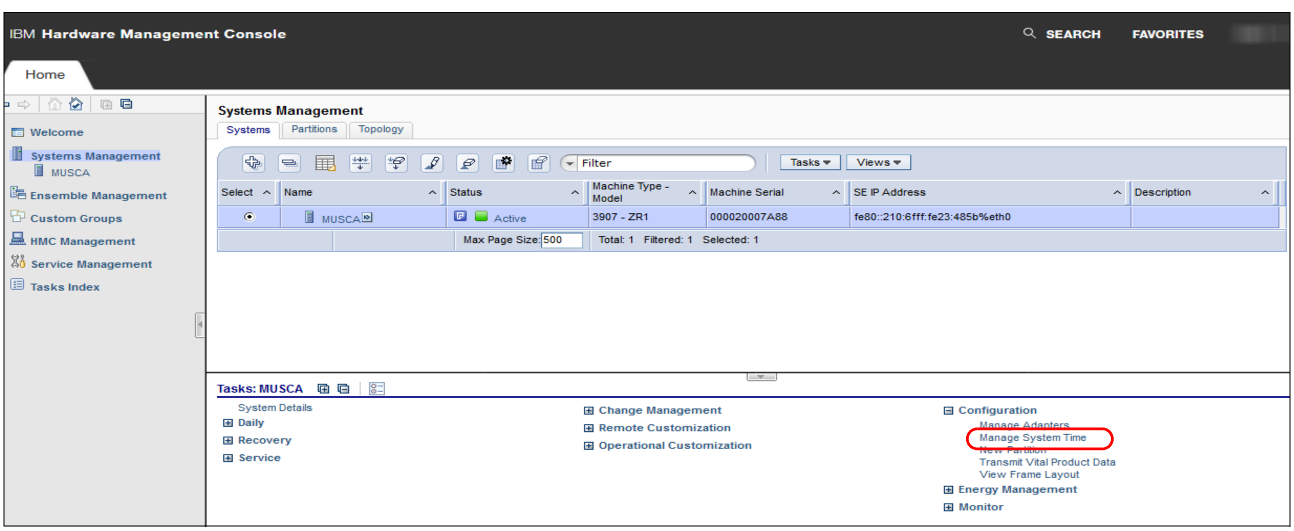

Figure 16-9 The Manage System Time task on the HMC

2. In our environment, the HMC is used to manage multiple CPCs. The Manage System Time task shows information about the CTN configuration for all configured CTNs.

   For a new CPC, no predefined CTN information is available; therefore, we enter a name (see HMC help) for our CTN.

   We use a created CTN with the ID `ITSO2`. The CPC we configured for DPM is not part of any CTN now. It is shown as STP stratum 0; therefore, it is considered a candidate for the `ITSO2` CTN, as shown in Figure 16-10.

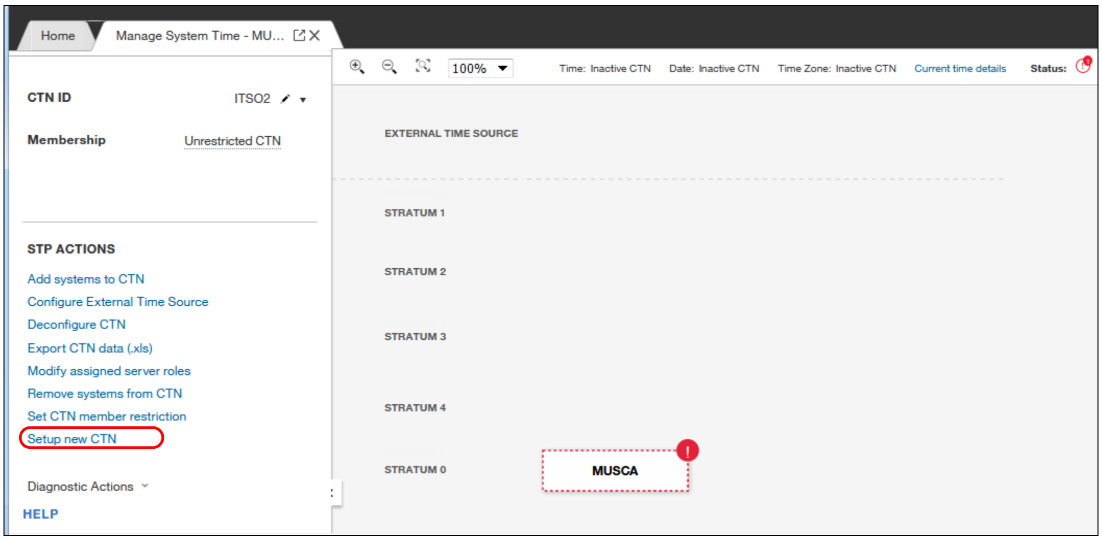

Figure 16-10 CTN candidate

3. We click **Setup new CTN** and start configuring the single server CTN, as shown in Figure 16-11. Because this server is a single server CTN, no Backup Time Server or Arbiter are available for this configuration.

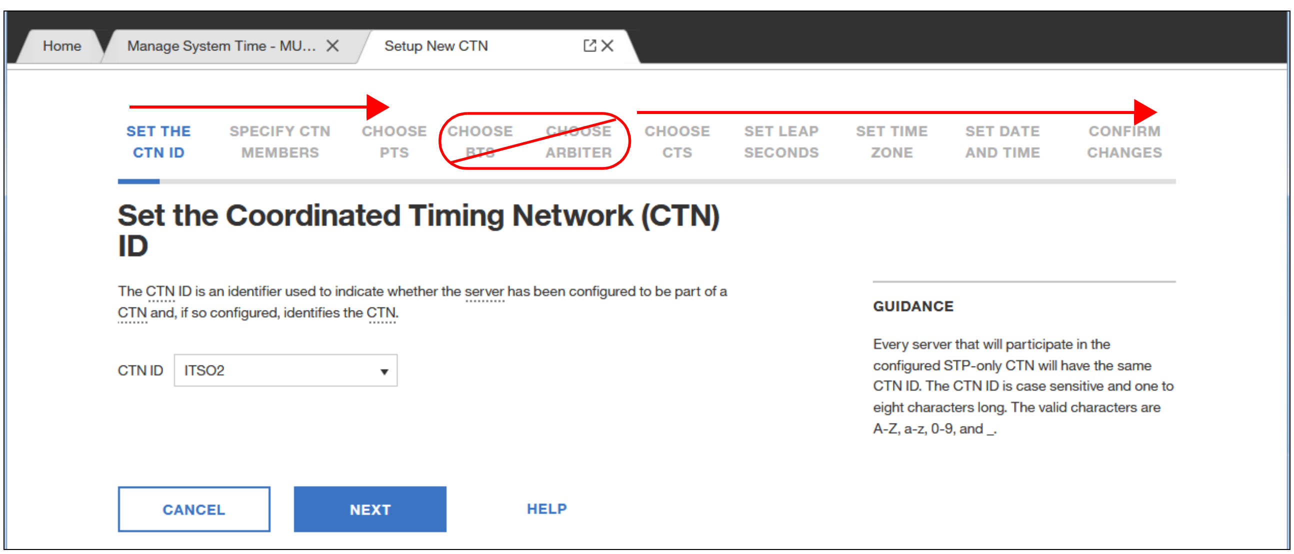

Figure 16-11 Setting up a new CTN

4. In our example, we specify our CPC as a member in the CTN, as shown in Figure 16-12 on page 398. Click **NEXT**.

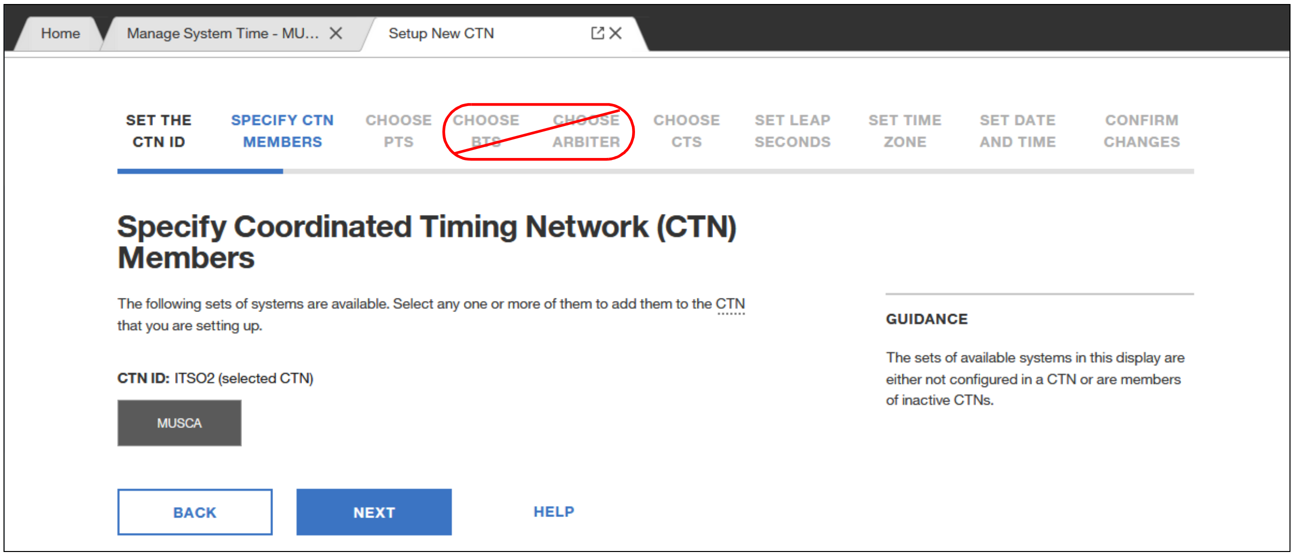

Figure 16-12 Choosing CTN member

5. We select our CPC (see Figure 16-13) as the Preferred Time Server (PTS) and skip selecting a BTS and Arbiter.

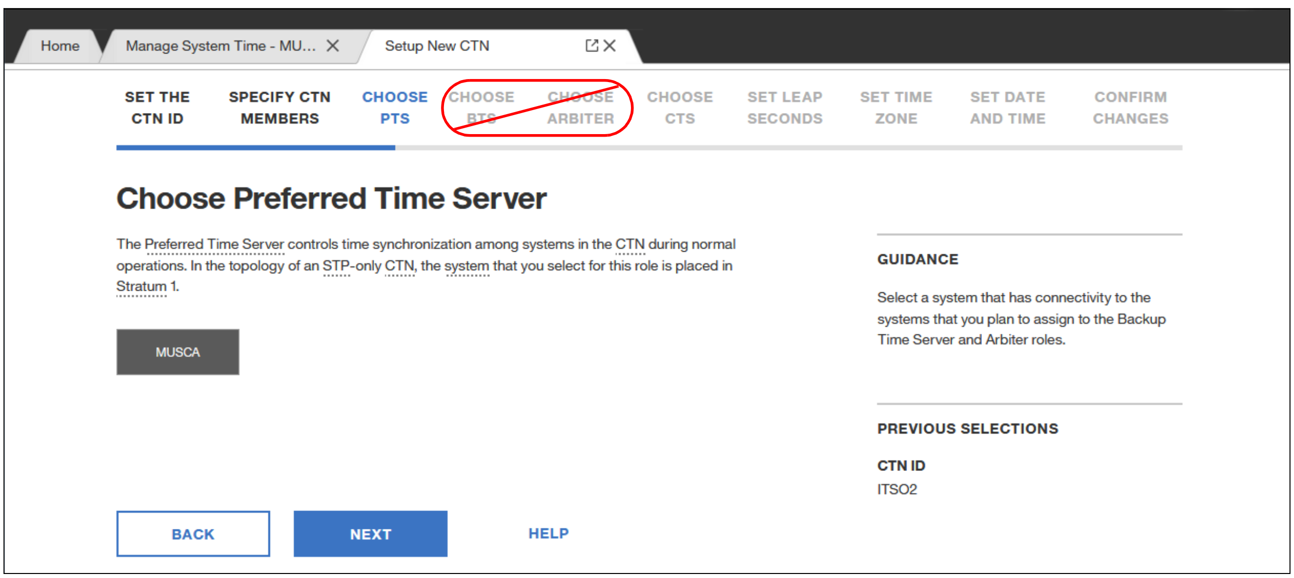

Figure 16-13 Selecting the PTS

6. We select the CTS as the Preferred Time Server (PTS), as shown in Figure 16-14.

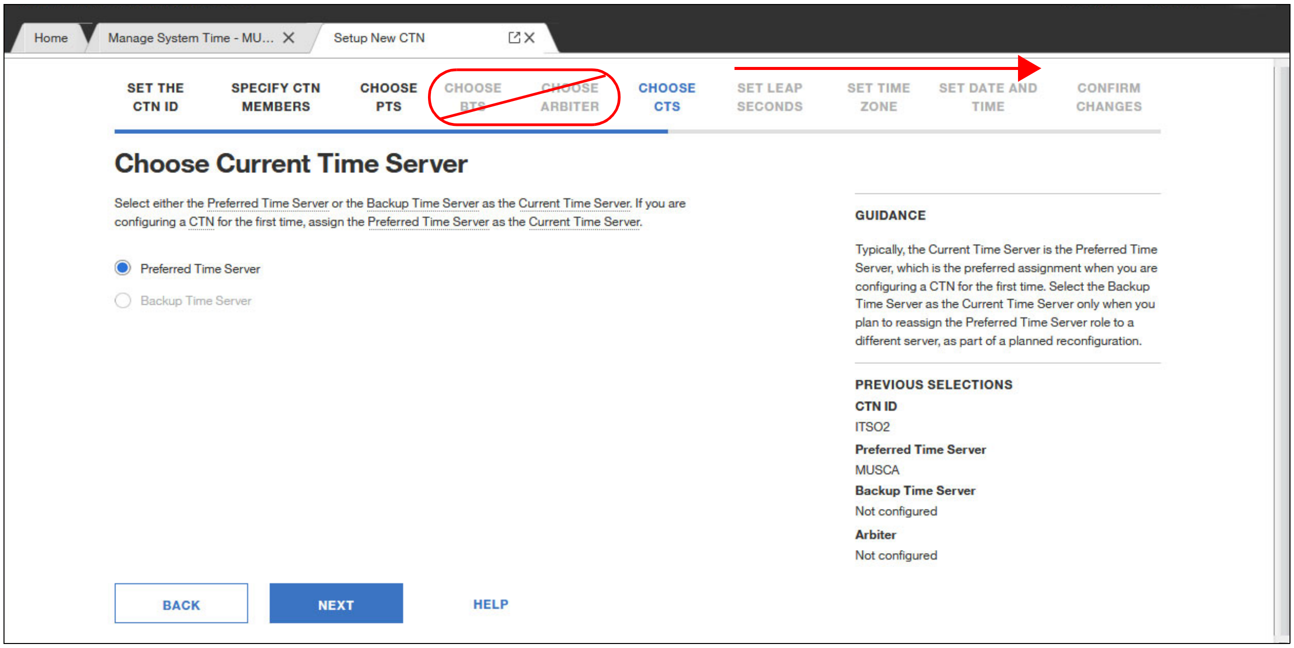

Figure 16-14 Choosing Current Time Server

7. We proceed to setting the Leap Seconds, Time Zone information, and we set the date and time, as shown in Figure 16-15. We configured the External Time Source (ETS) in steps 1 - 4, as described in 16.2.1, “Configuring External Time Source on the Support Element” on page 394.

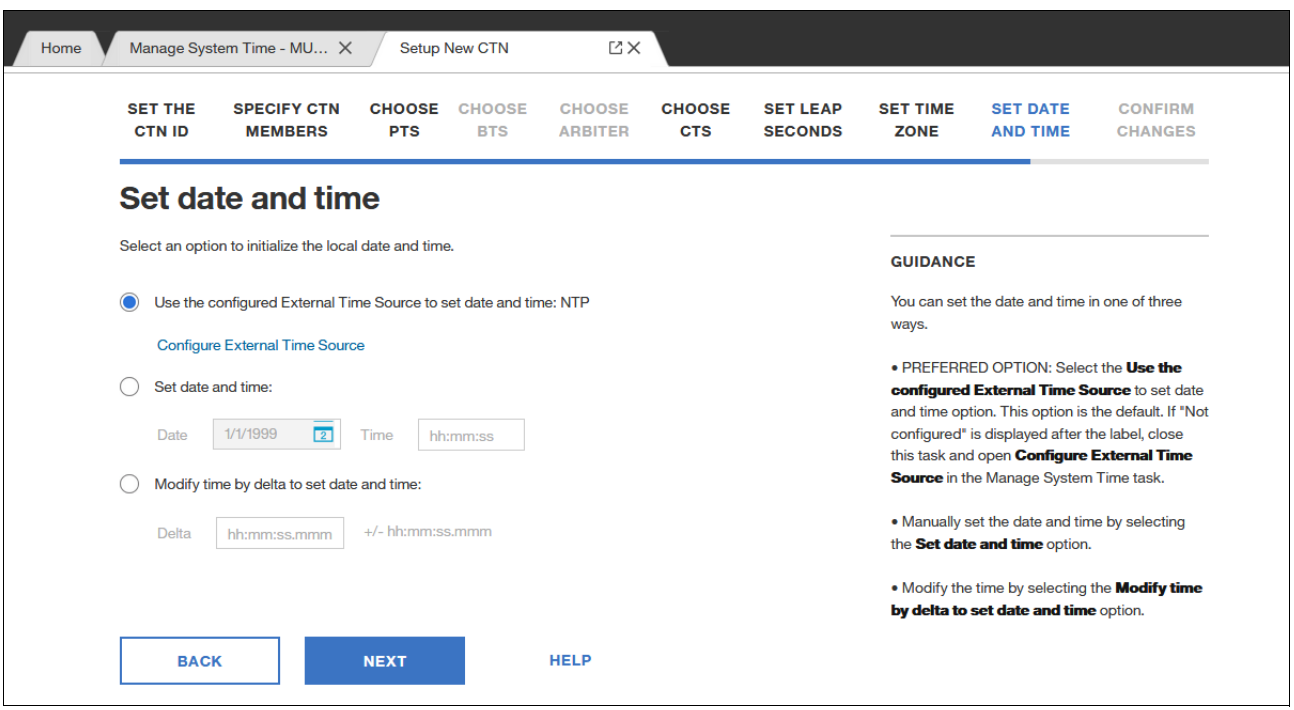

Figure 16-15 Set date and time

8. After clicking **NEXT**, we check the configuration before confirming any changes, as shown in Figure 16-16.

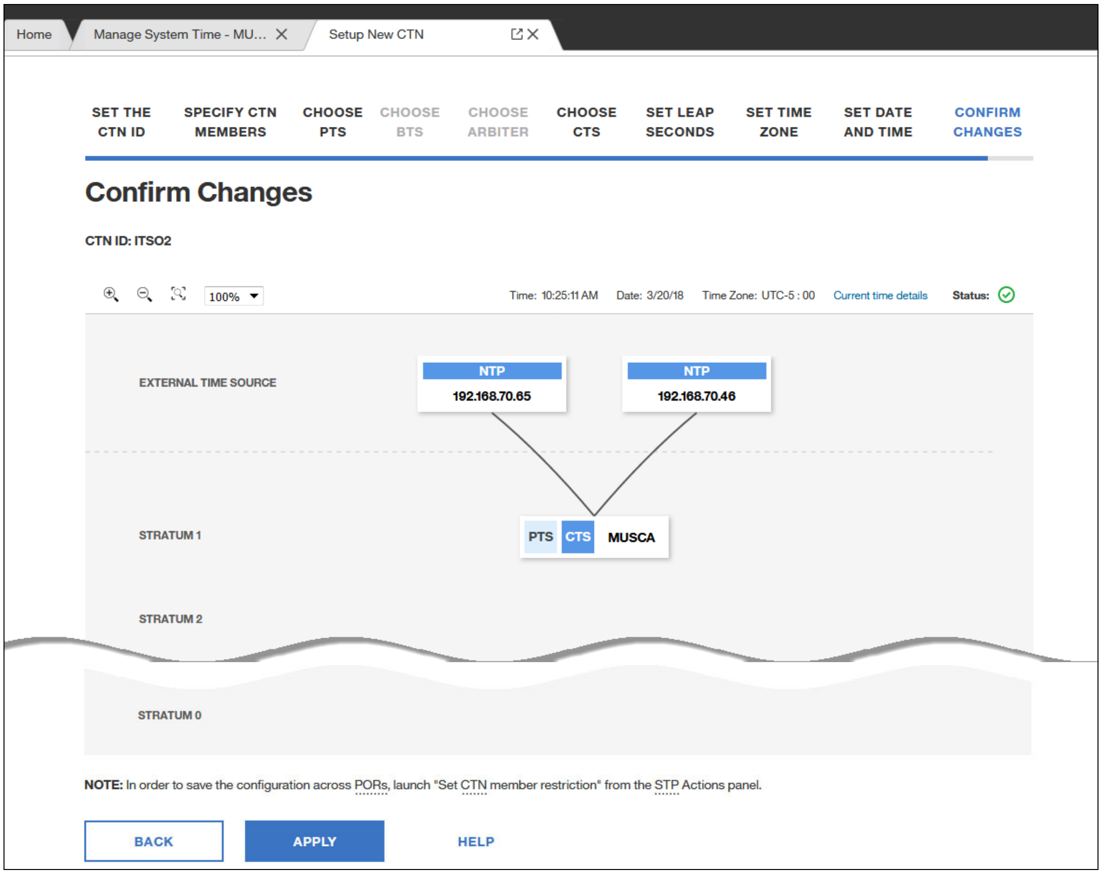

Figure 16-16 Single Server CTN (pending Confirm Changes)

9. After the configuration is applied, the CTN is active and the CPC time of day is directed to the ETS. The result is shown in Figure 16-17 on page 401.

Saving the CTN configuration across PORs: The resulting default configuration allows other CPC to be added to this CTN ("unrestricted CTN"). However, because no coupling links are used, this addition is not possible.

Also, the configuration of an unrestricted CTN does not persist after a Power-on-Reset (POR); that is, the configuration is not saved. Therefore, the CTN must be reconfigured after a POR.

To save the CTN configuration across PORs, the CTN membership must be restricted as described next.

T

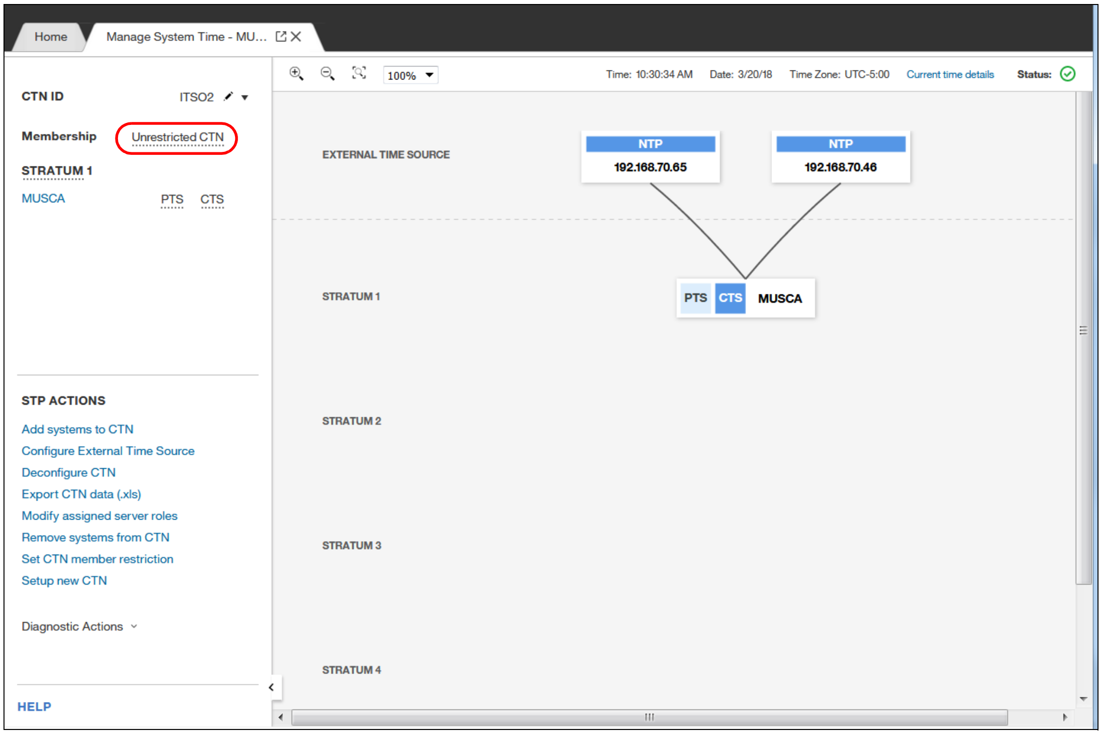

Figure 16-17 CTN configuration result - Unrestricted CTN

10. We restrict the CTN membership to save (persist) the configuration after a PRO, as shown in Figure 16-18.

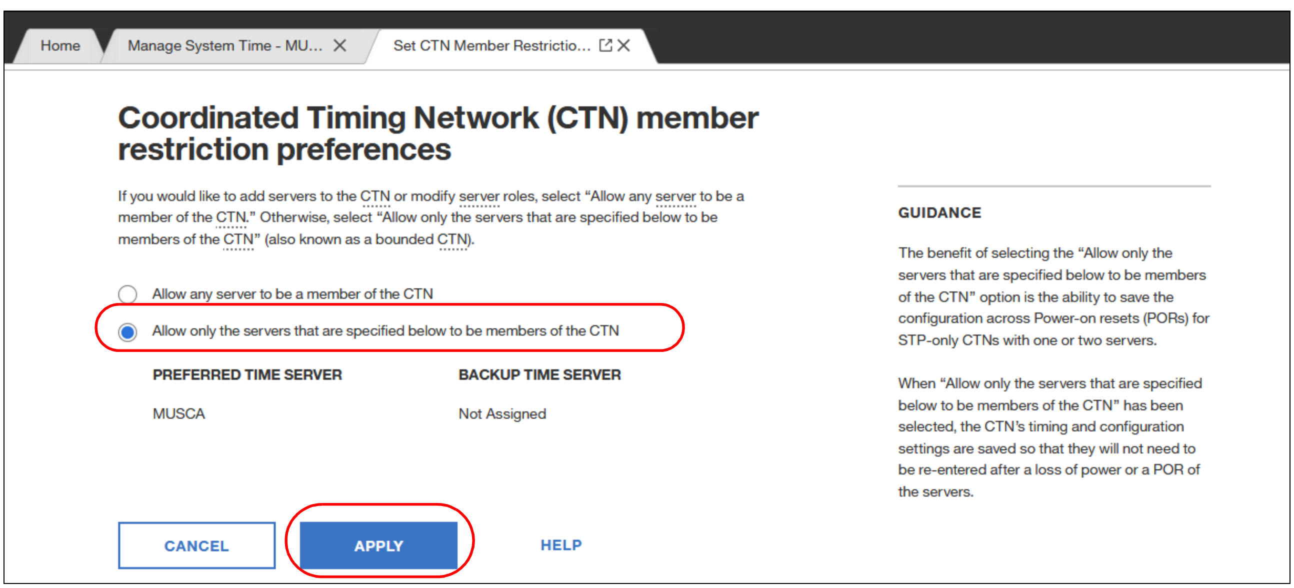

Figure 16-18 Restricting CTN membership

16.3 Configuring partitions

In this section, we describe how to configure three types of supported partitions: Linux, z/VM, and SSC.

After the CPC is activated in DPM mode, you can configure the system. The configuration process includes the following high-level tasks:

- Planning and checking available resources (processor, memory, adapters, accelerators, cryptographic cards, and devices)
- Configuring storage adapters and environment
- Configuring partitions

For more information about how to configure your system by using DPM, see *IBM Dynamic Partition Manager (DPM) Guide*, SB10-7170, which is available at the IBM Resource Link website (login required).

16.3.1 Checking system resources

This section describes how to browse the interface for more information about system resources.

System overview

Figure 16-19 shows a system overview of CPC that is managed by DPM.

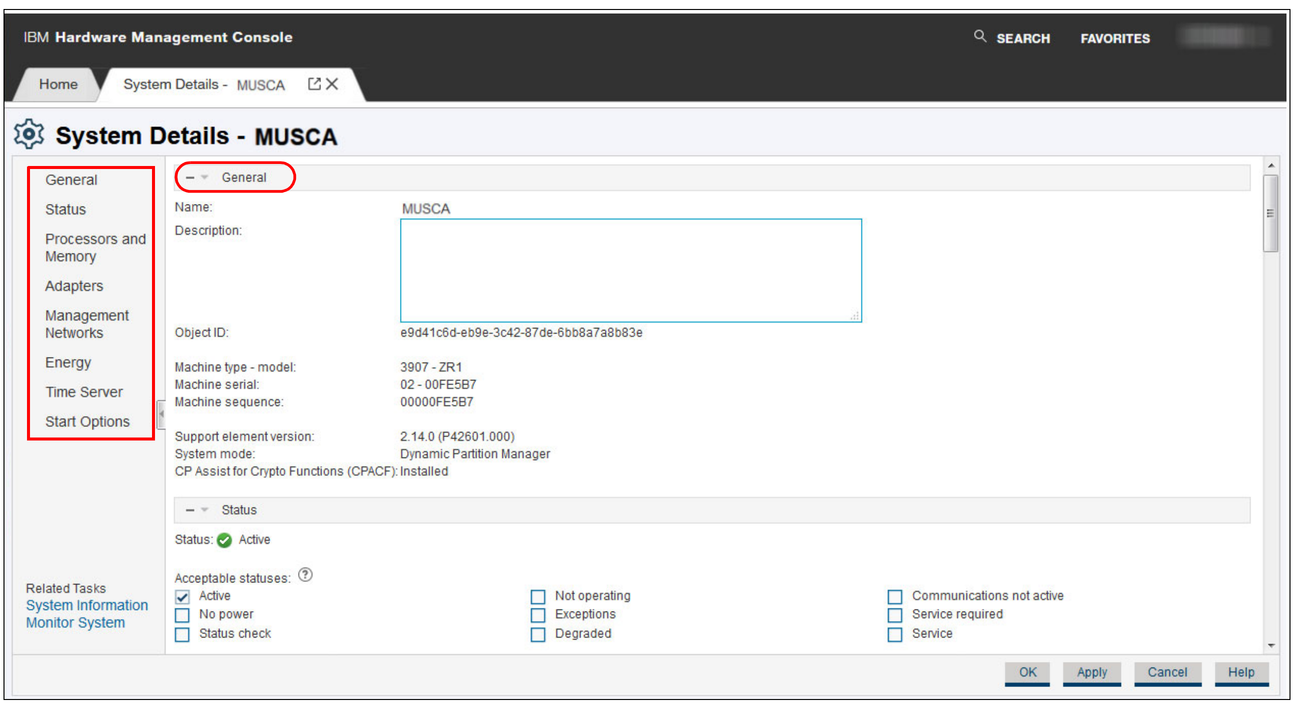

Figure 16-19 System Details: General information

Figure 16-20 shows the system status.

Figure 16-20 Status window

Processor and Memory

The installed and available processor and memory resources are shown in Figure 16-21.

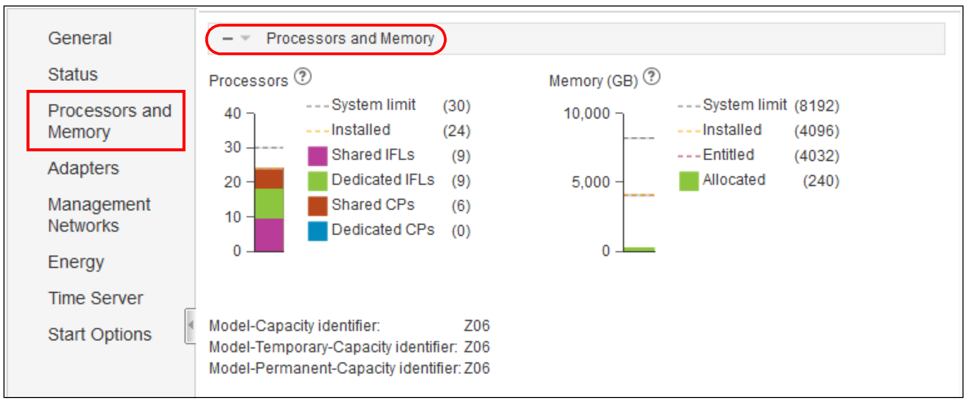

Figure 16-21 Processors and Memory window

Adapters

The installed adapters summary usage information is shown in Figure 16-22.

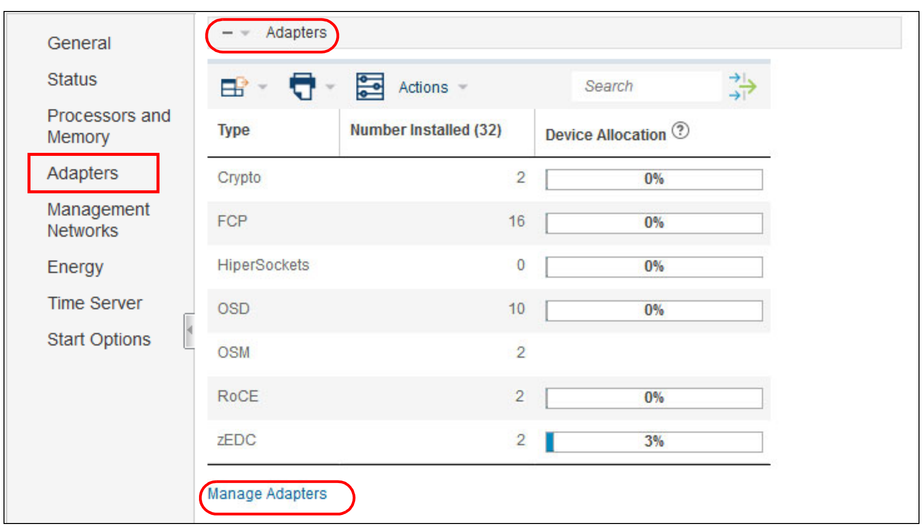

Figure 16-22 Adapters summary window

Adapters information

To check the detailed list of adapters, click **Manage Adapters** (see Figure 16-22 on page 403). A new tab opens in which detailed adapter information is listed (see Figure 16-23). You can filter the adapters based on functionality, as highlighted at the top of the table.

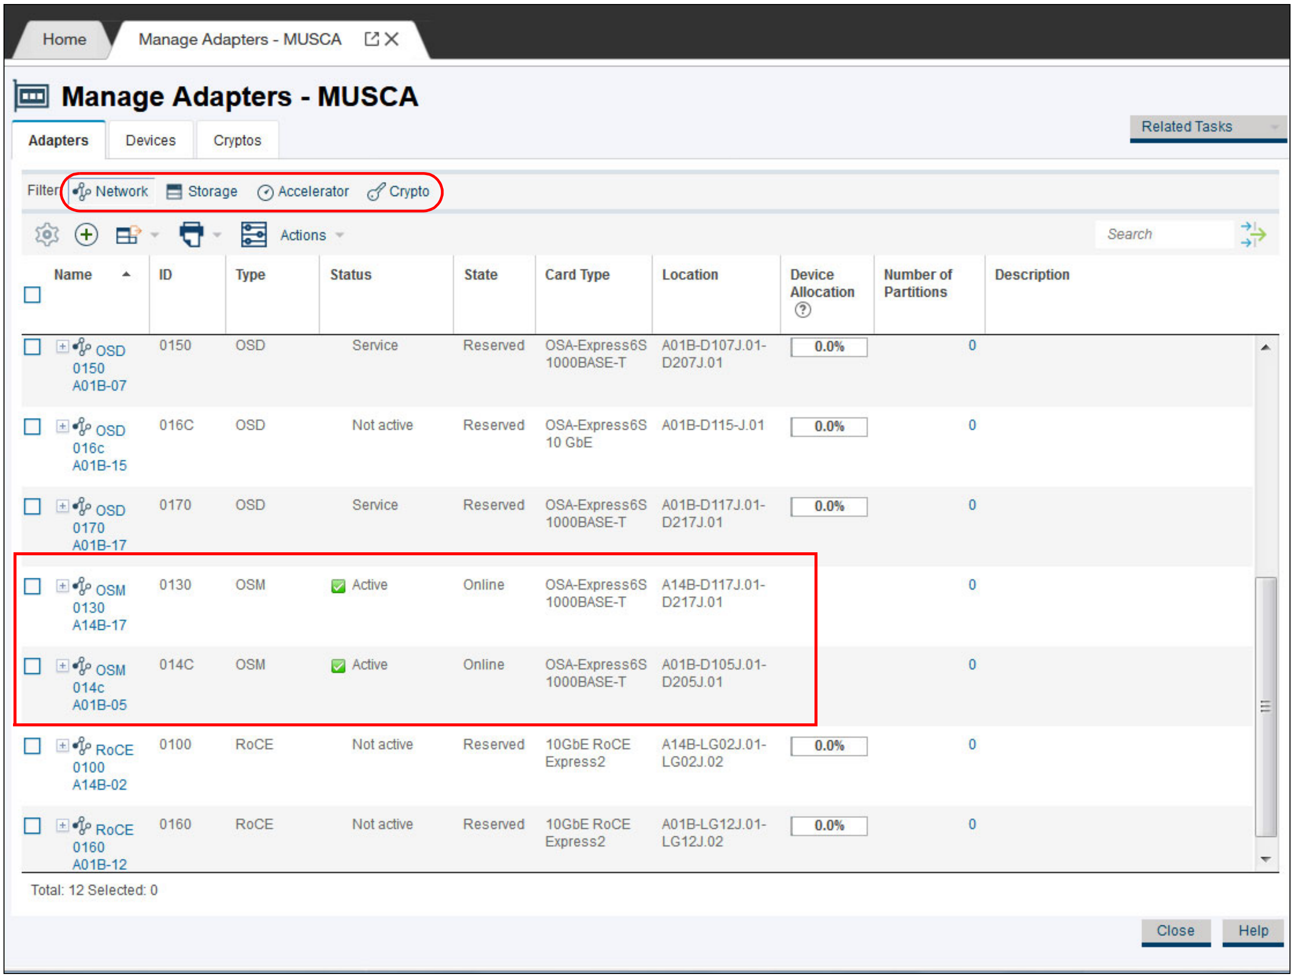

Figure 16-23 Adapter list

> **Tip:** The OSA-Express Adapters that are highlighted in the middle of Figure 16-23 are configured as OSM (management adapters). These two adapters cannot be used for any other purpose.

16.3.2 Preliminary task: Storage configuration

Starting with DPM 3.1, support for FICON (ECKD storage) and FCP (SCSI storage) is provided. DPM provides a Storage configuration task, which is required before allocating external storage to your LPARs.

You can filter the adapter list to identify the available storage adapters, as shown in Figure 16-24 on page 405. The adapters are shown as “Unconfigured”.

Now, you must proceed with the storage configuration to select the protocol to be used for storage access for each adapter and configure the storage (access the storage subsystem, allocate storage space, and so on).

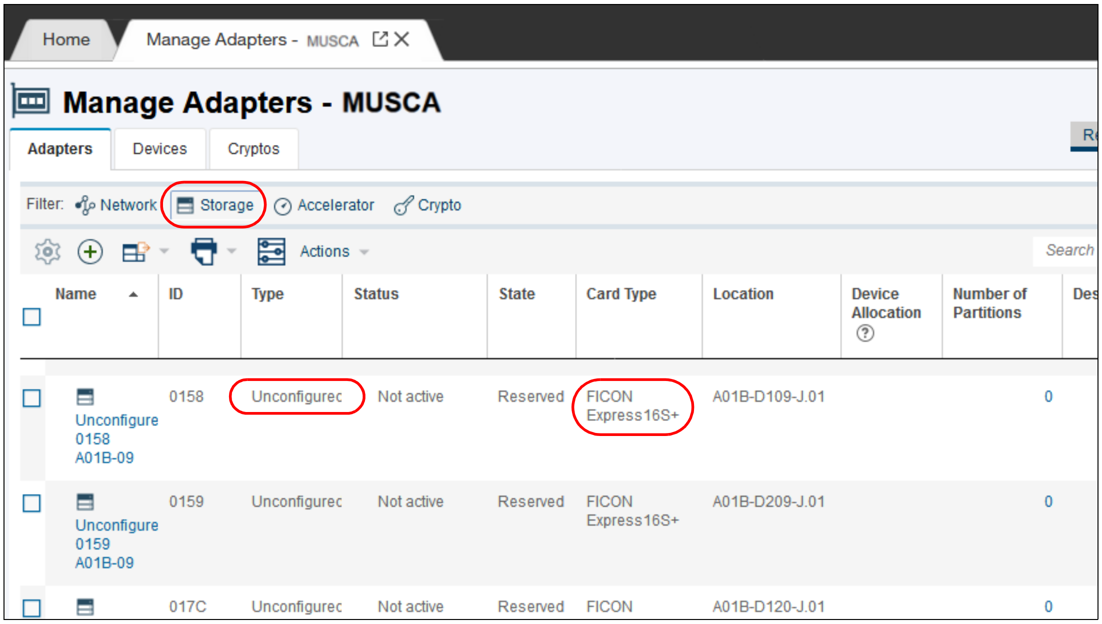

Figure 16-24 Checking available FICON Express adapters

Complete the following steps:

1. Click **Configure Storage** (see Figure 16-25) to access the storage configuration menu.

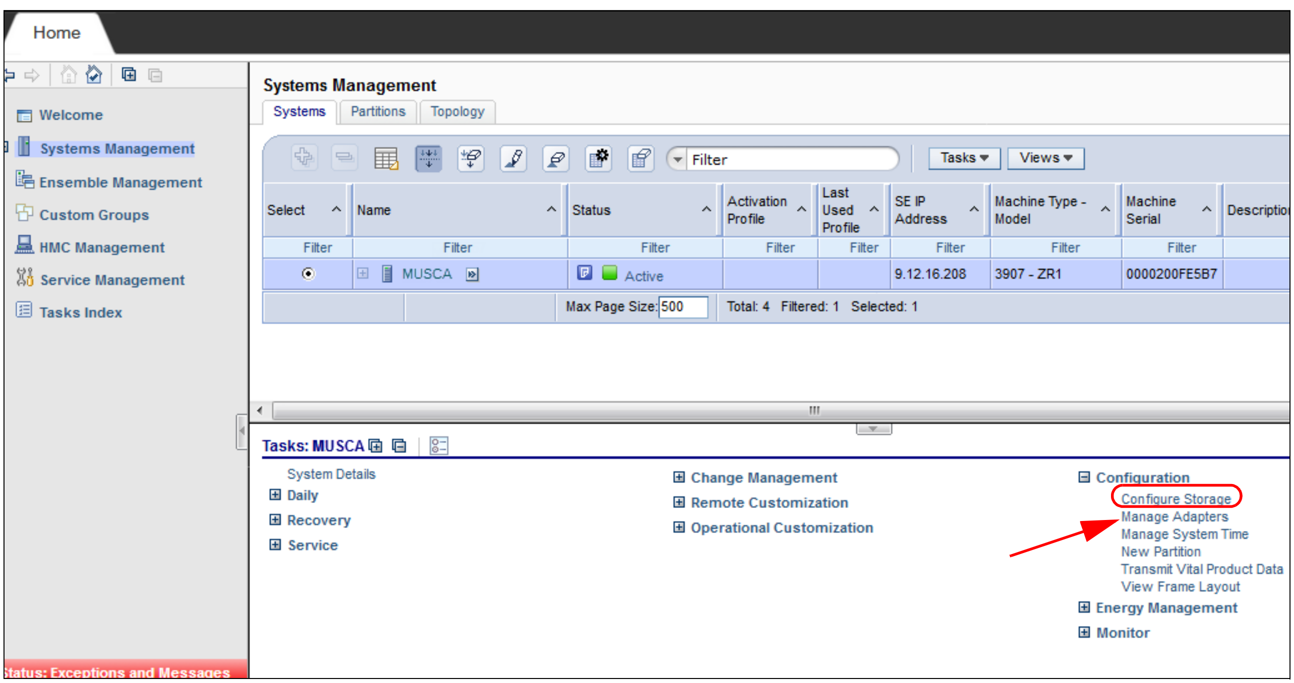

Figure 16-25 Accessing the "Configure Storage" task

2. In the Configure Storage window, configure the cards' (adapters') protocol (FICON or FCP), request storage, and configure FICON connections (switches/directors and storage devices).

In our test environment, we used FCP (SCSI) storage; therefore, we selected **FCP** for our adapters, as shown in Figure 16-26.

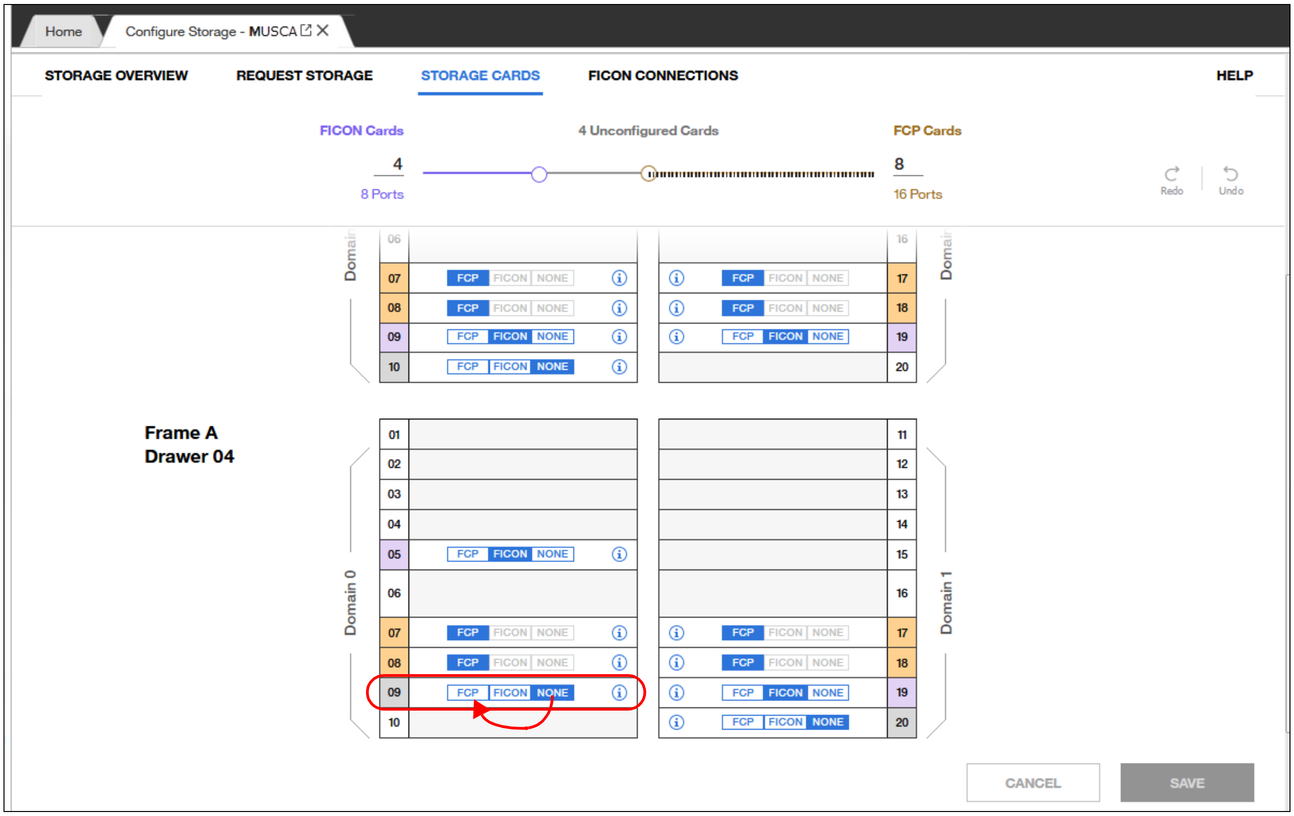

Figure 16-26 Storage cards configuration

Identify the adapters that you plan to use (unused adapters are marked as NONE), click the wanted protocol (FCP or FICON). Then, save the configuration and return to main menu (see Figure 16-25 on page 405).

3. After the storage adapter protocol is saved, the status of the adapter is reflected in the Manage Adapters task (see the red arrow Figure 16-25 on page 405), as shown in Figure 16-27 on page 407.

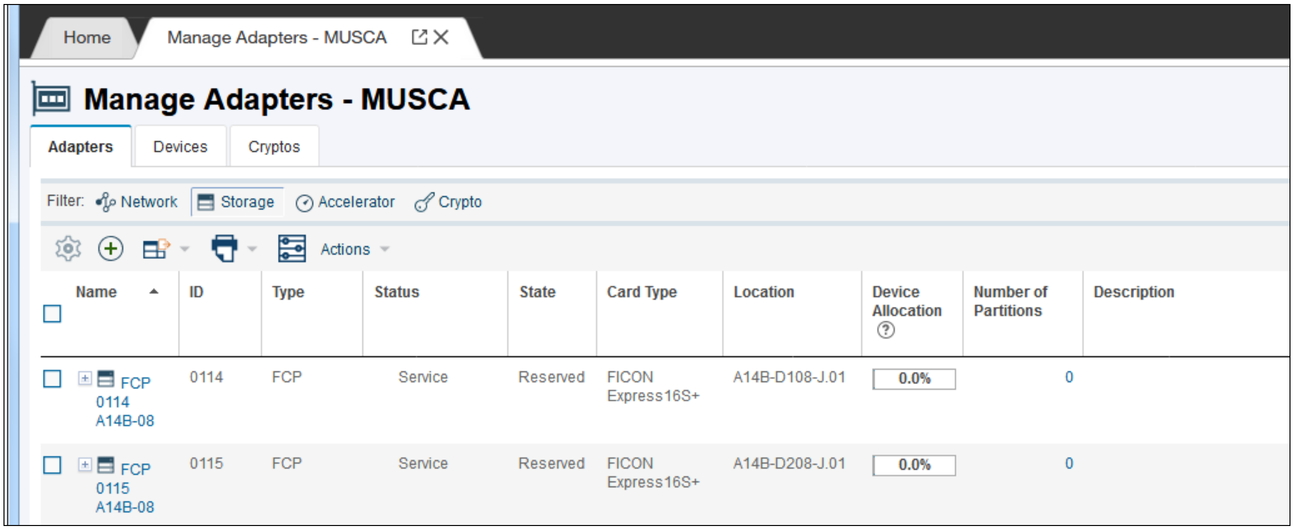

Figure 16-27 Storage adapters available for allocation to partitions

For more information about configuring FICON (ECKD) storage, see *IBM Dynamic Partition Manager (DPM) Guide*, SB10-7170.

16.3.3 Configuring a Linux partition

This section describes how to create a Linux partition by using the Standard configuration menu. We present the same steps for configuring a partition by using the Advanced menu as described in 16.3.6, “Configuring partitions by using Advanced menus” on page 424.

We used the following process:

1. To configure a partition, we start with the Systems Management interface on the HMC, select the CPC to be configured. Then, in the Tasks pane, we select **New Partition**, as shown in Figure 16-28.

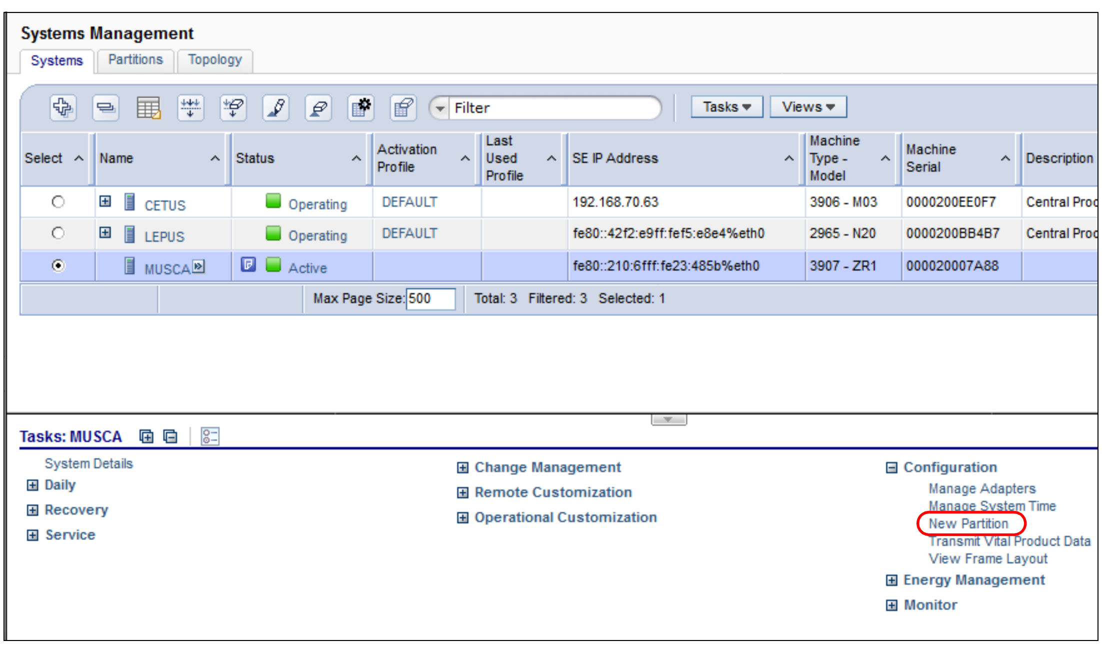

Figure 16-28 New Partition task selection

2. The menu that is shown in Figure 16-29 opens. We select **Okay, got it** to use the standard partition configuration menus.

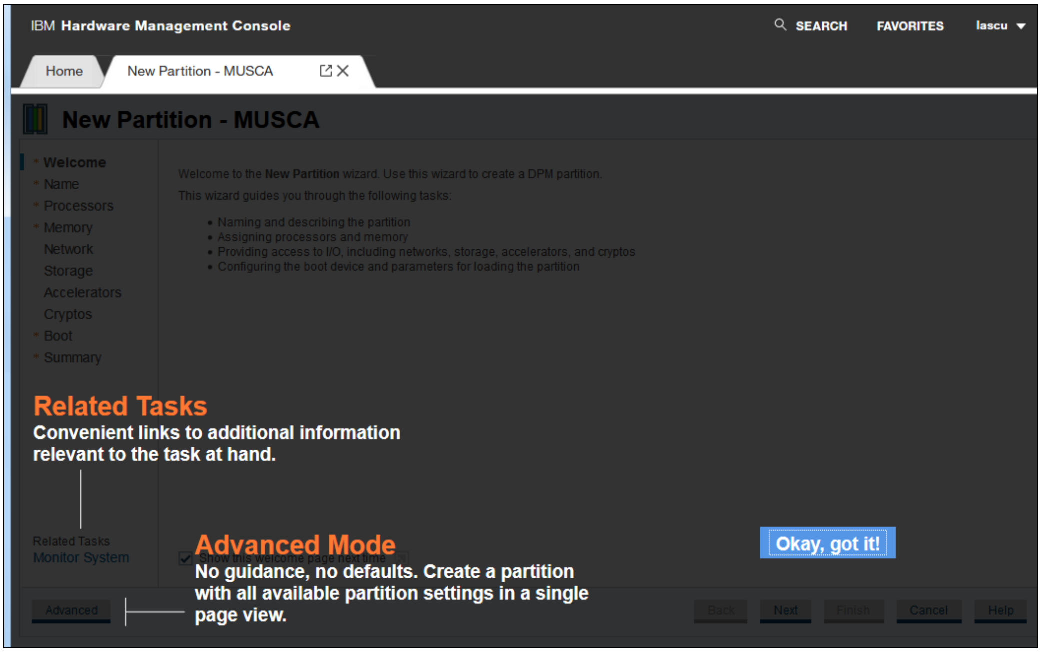

Figure 16-29 Emphasizing the configuration menus options (Standard / Advanced)

Figure 16-30 shows the standard New Partition configuration (wizard) menu. Later in this chapter, we also create an LPAR by using the Advanced menus.

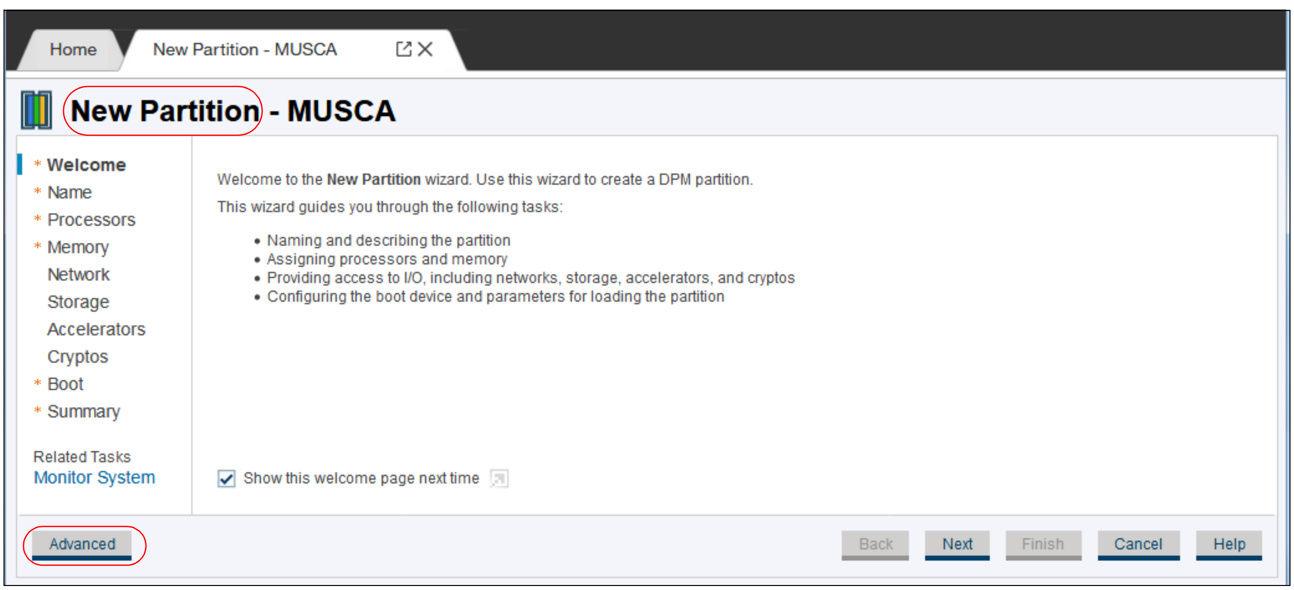

Figure 16-30 Initial menu for configuring partitions (New Partition wizard)

3. The basic information for a partition (name, description, and type) are configured in the next menu, as shown in Figure 16-31. We select the partition type as **Linux** and click **Next**.

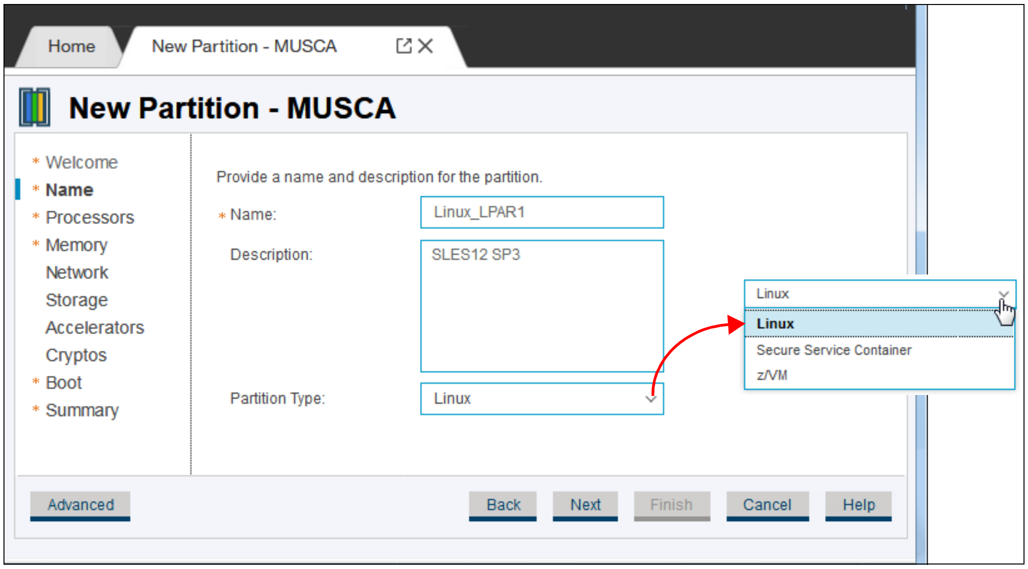

Figure 16-31 Partition name, description, and type

Depending on the partition type, several processor types (as they are characterized in the machine configuration) can be used. For Linux partitions, only CP or IFL processor types are supported, as shown in Figure 16-32.

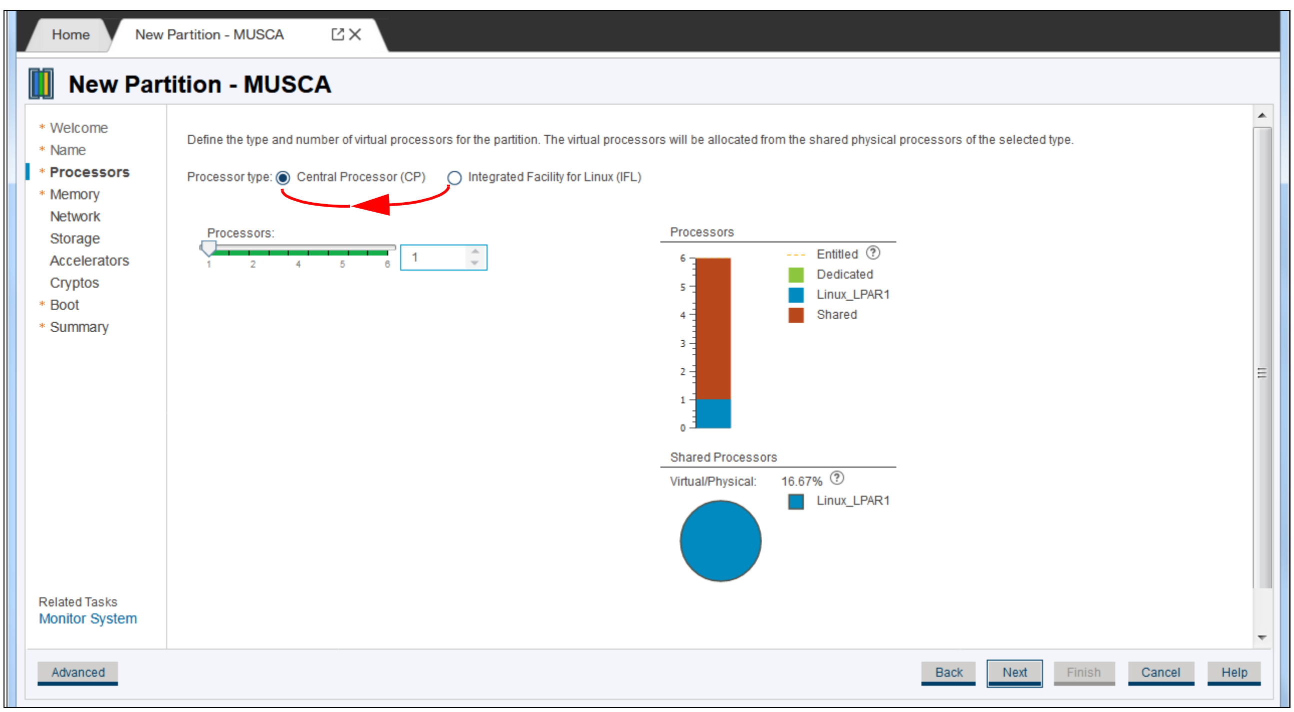

Figure 16-32 Available processor types for Linux partition

4. We select for our partition processor type IFL, as shown in Figure 16-33. Then, click **Next**.

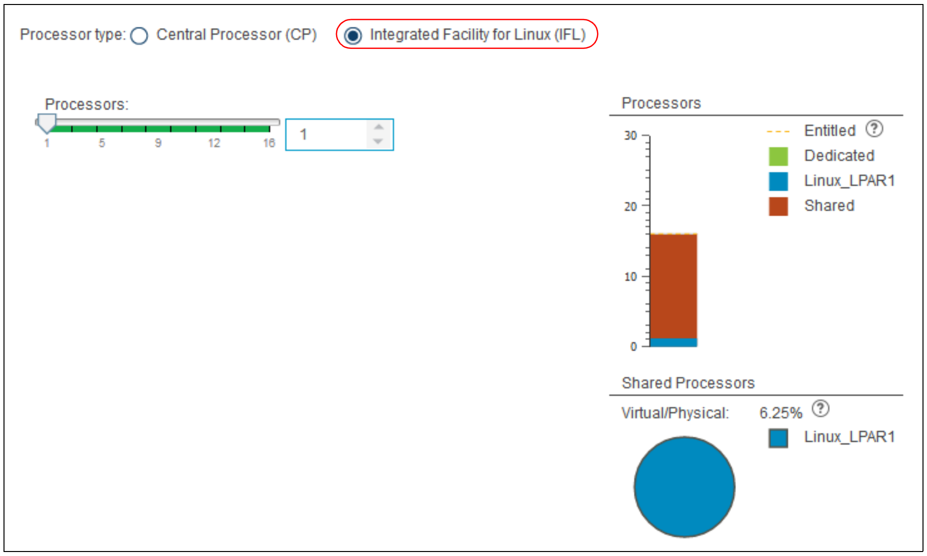

Figure 16-33 Selecting IFLs for the partition

5. In the next menu (see Figure 16-34), we select the amount of memory to be made available for our partition. The "desired" amount is 16 GB and the maximum amount is 24 GB.

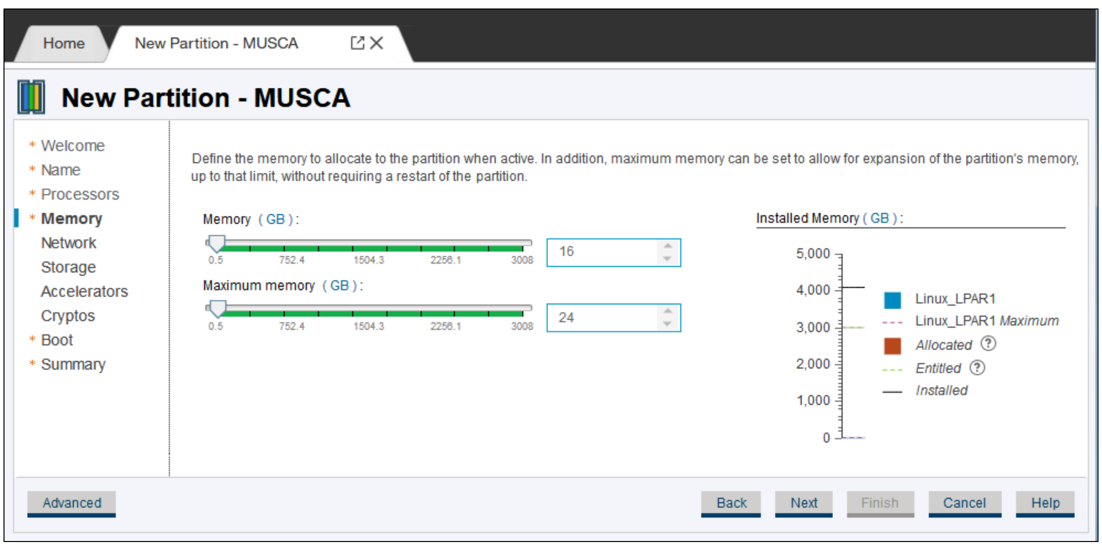

Figure 16-34 Memory selection

6. After the memory is selection, we click **Next** and start configuring adapters for our partition. Figure 16-35 on page 411 shows the Network Interface Card (NIC) selection window.

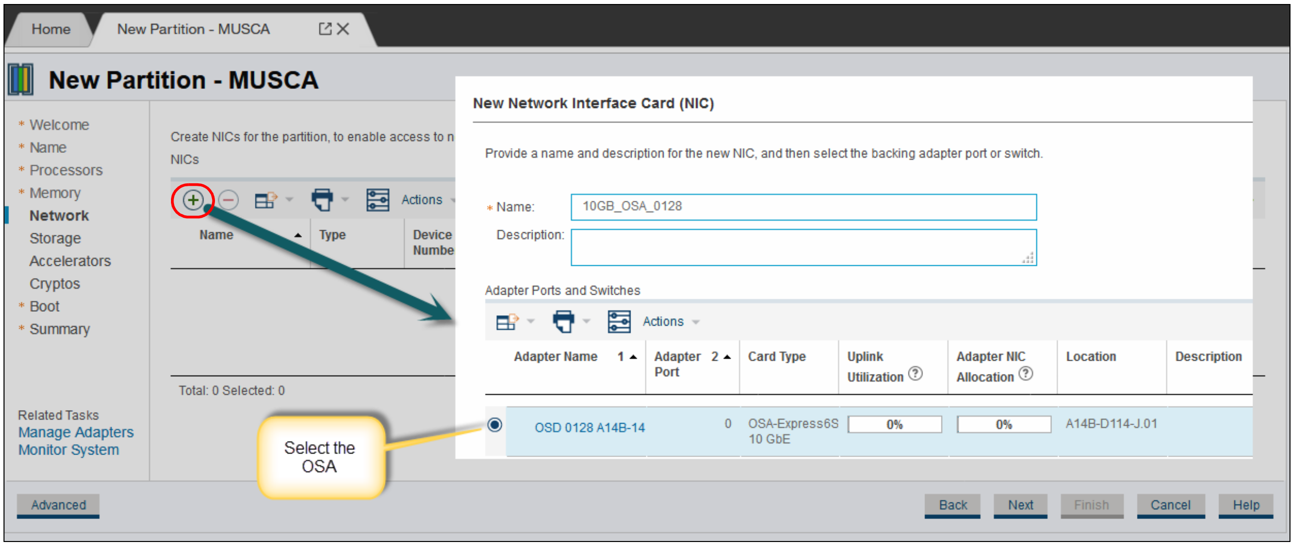

Figure 16-35 NIC selection

7. We selected an OSA-Express adapter and created one device (Device number 0001) for our partition, as shown in Figure 16-36.

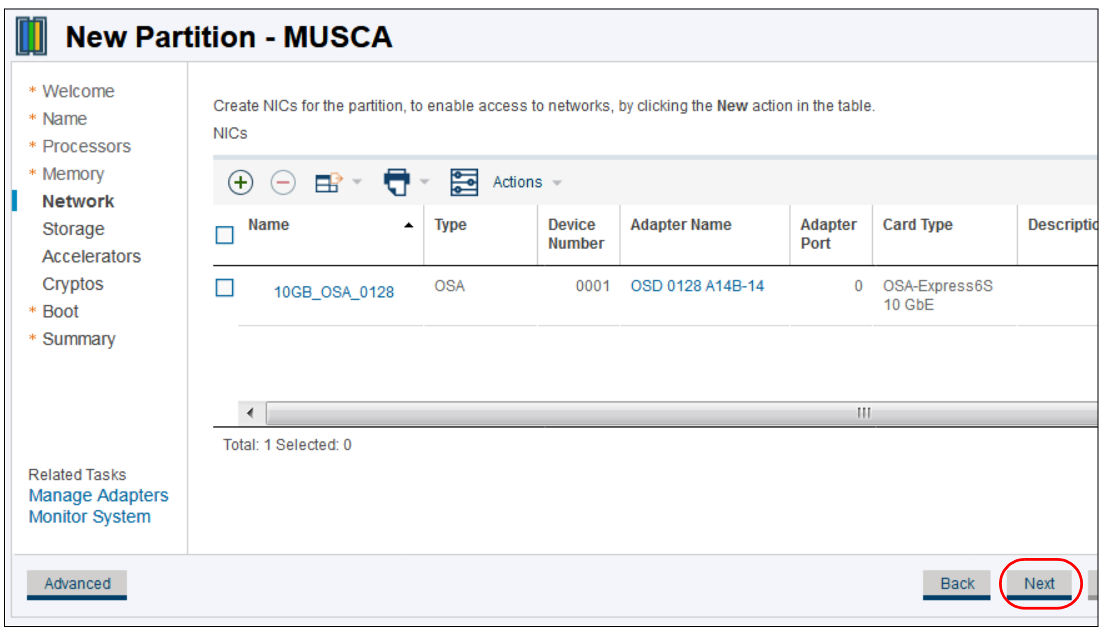

Figure 16-36 Network device configured

8. We select a storage access (host bus) adapter for our partition, as shown Figure 16-37 on page 412. Because we use SCSI protocol for storage access, we select an adapter that is configured as FCP.

Note: For more information about Linux on Z device support and configuration, see the IBM Knowledge Center page for Linux on Z and LinuxONE.

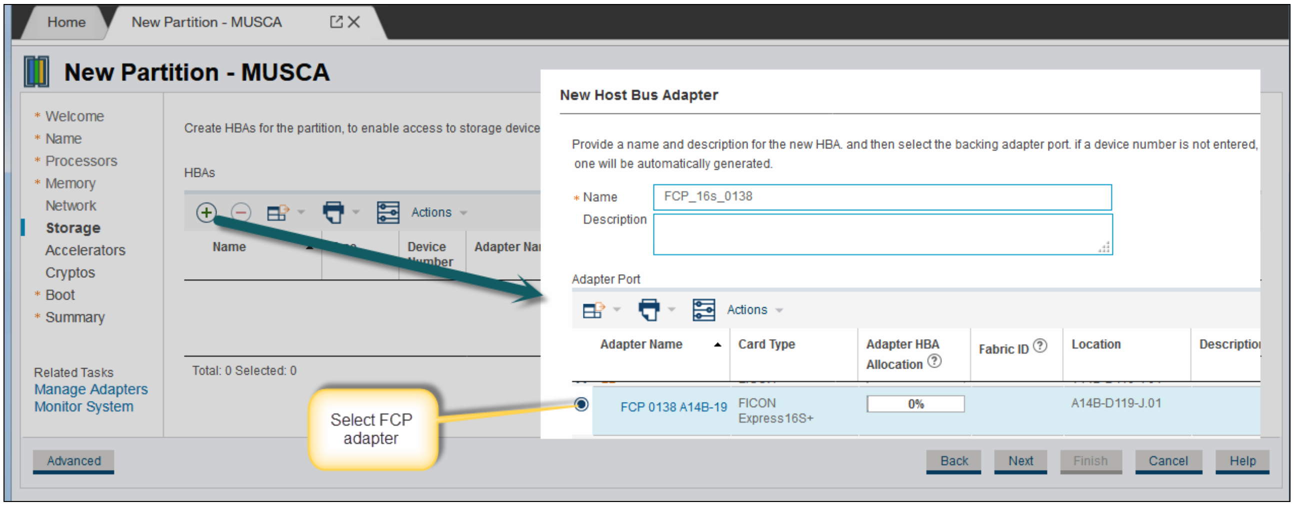

Figure 16-37 Selecting FCP adapter

Figure 16-38 shows the NPIV device (device 0004) that was created for our partition.

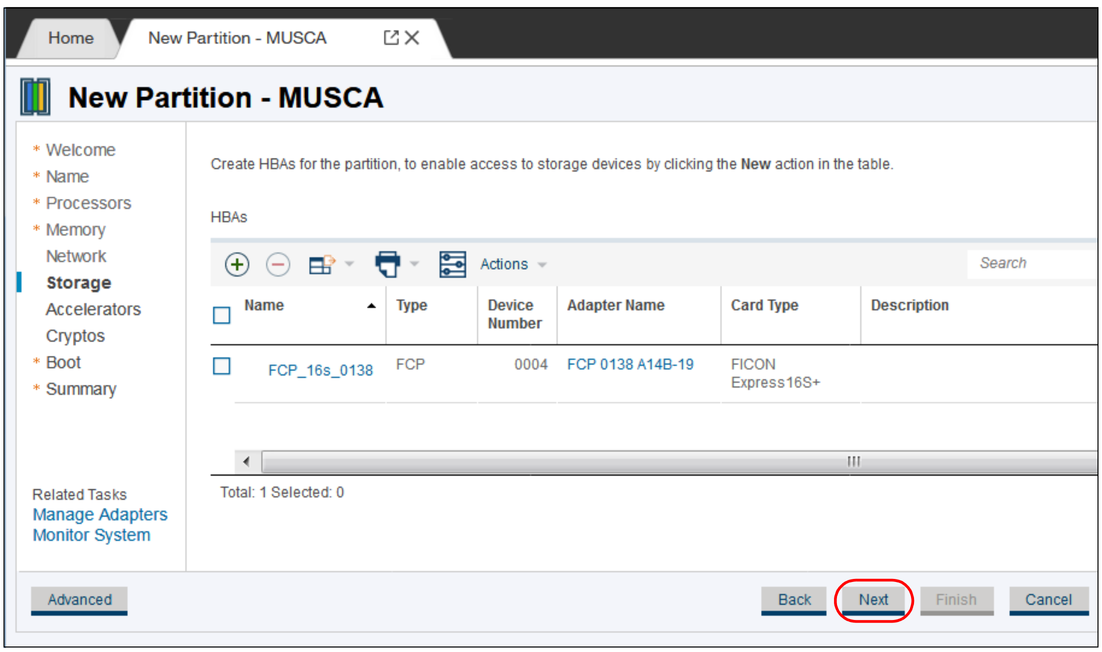

Figure 16-38 FCP (NPIV) device

9. If wanted, we select Accelerators and Cryptos. However, these components can be added later. We chose to skip Accelerators and Cryptos configuration and proceed to selecting Boot options, as shown in Figure 16-39.

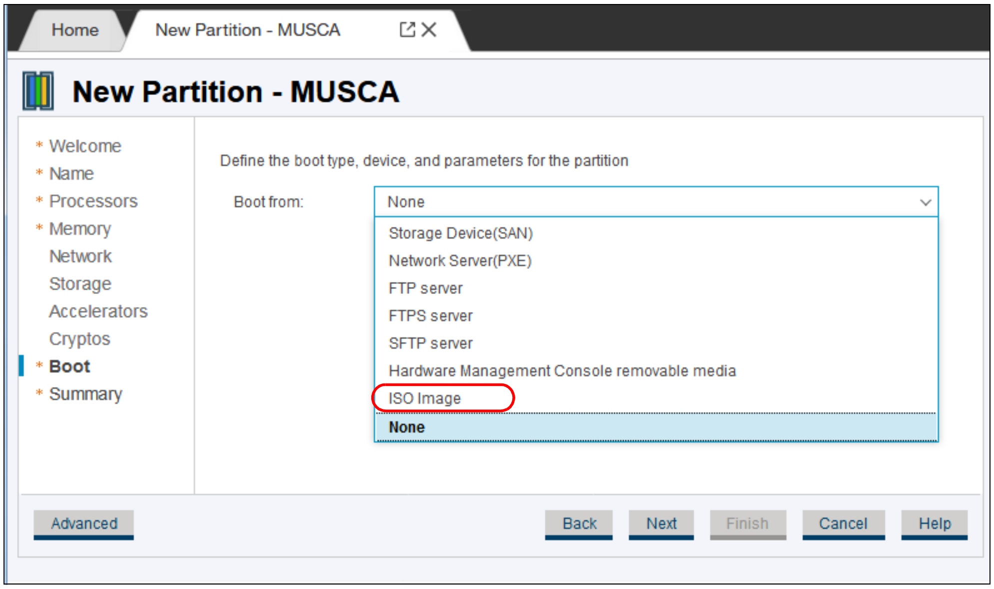

Figure 16-39 Selecting Linux partition boot options

10. After selecting the partition boot options, we finalize creating the partition by clicking **Finish**, as shown in see Figure 16-40.

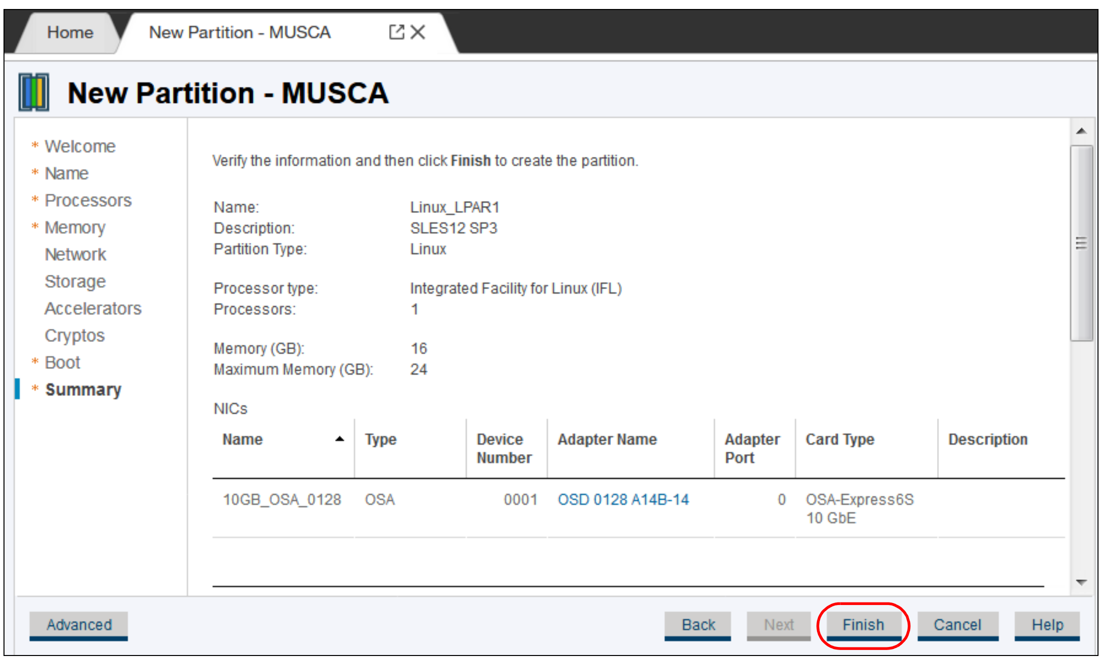

Figure 16-40 Finalizing partition definition

The options that are available when the partition is created are shown in Figure 16-41.

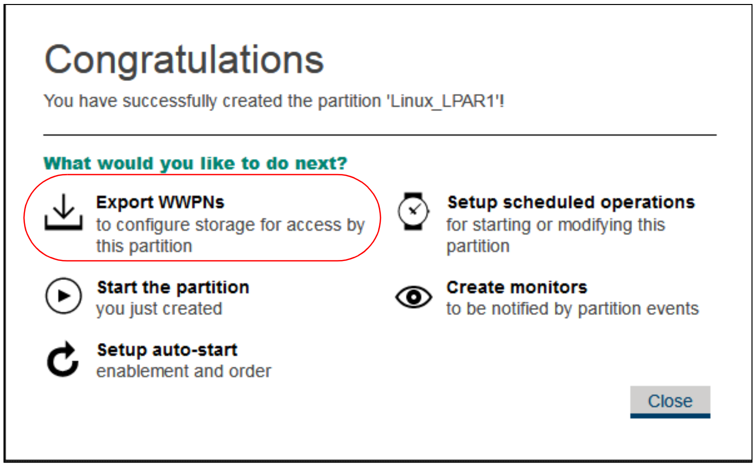

Figure 16-41 More task options after partition creation

11. Because we need information about the partition's storage device configuration, we chose to export WWPNs. WWPNs information is required for SAN zoning and storage access (LUN Masking).

Note: This book does *not* describe how to install an operating system or any other software in a partition. For more information about such installations, see your software of choice (operating system, hypervisor, or appliance) installation documentation.

16.3.4 Configuring a z/VM partition

This section describes configuring a partition for running the z/VM hypervisor.

Because the procedure is similar to configuring a Linux partition, we emphasize only the differences from the procedure that is described in 16.3.3, “Configuring a Linux partition” on page 407. For more information, see 16.3.6, “Configuring partitions by using Advanced menus” on page 424.

Complete the following steps:

1. To configure a partition, start with the Systems Management interface on the HMC. Select the CPC to be configured. Then, in the Tasks pane, select **New Partition**, as shown in Figure 16-42.

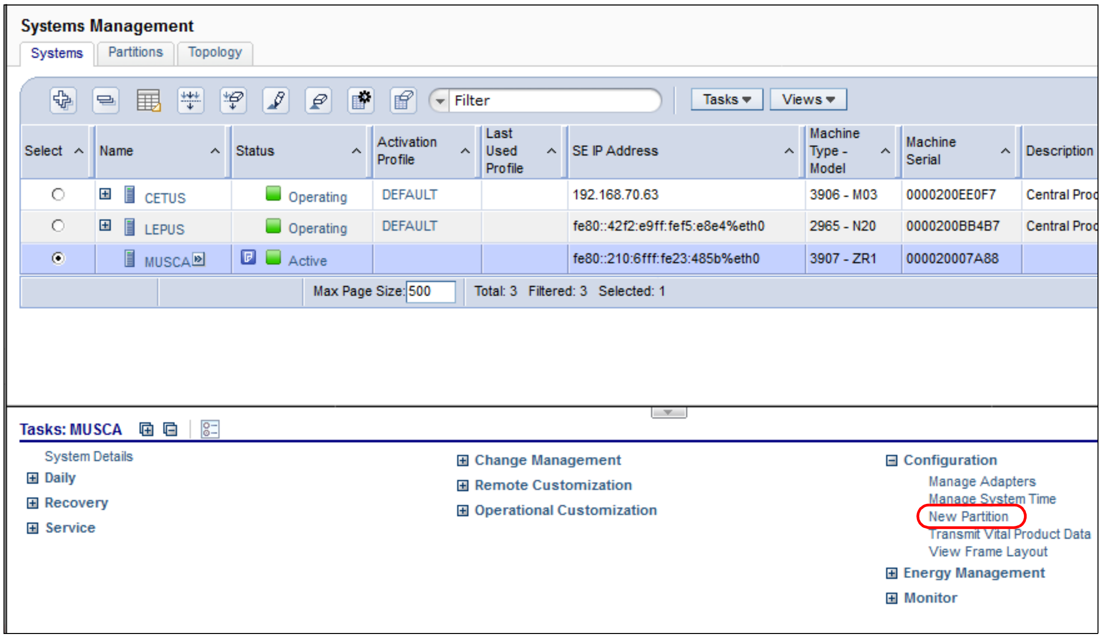

Figure 16-42 New Partition task selection

Figure 16-43 shows the standard New Partition configuration (wizard) menu. Later on in this chapter, we also create an LPAR by using the Advanced menus.

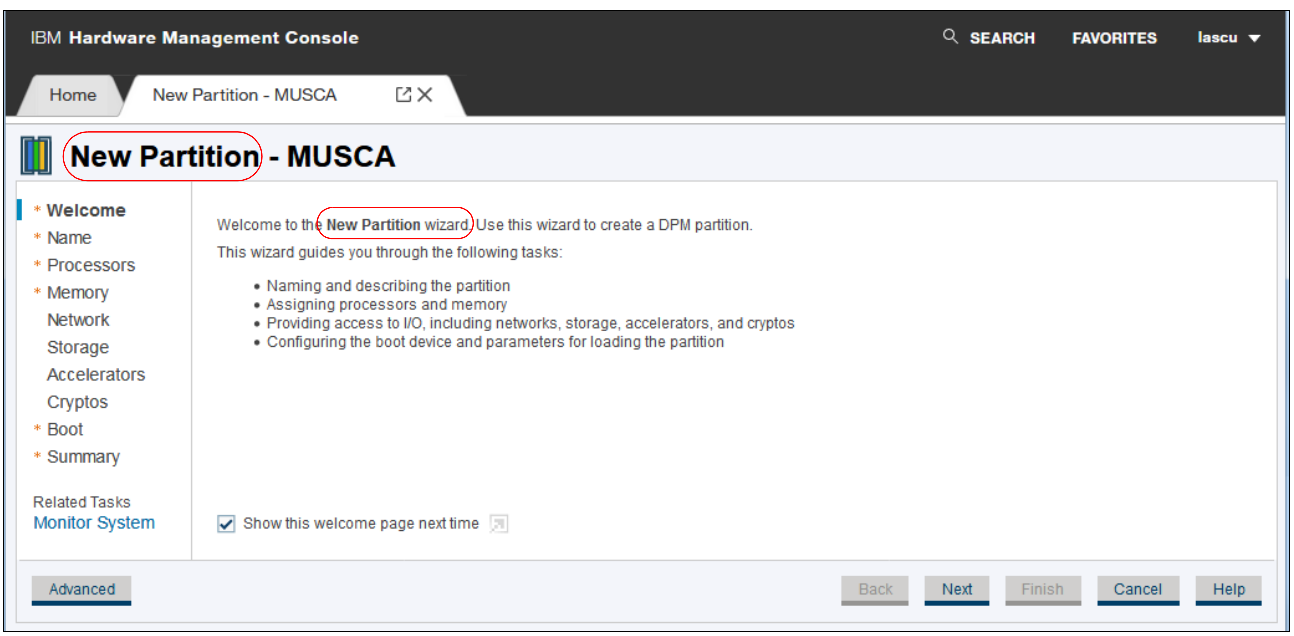

Figure 16-43 Initial menu for configuring partitions (New Partition wizard)

2. The basic information for a partition (name, description, and type) are configured in the next menu, as shown in Figure 16-44. Select the partition type as **z/VM** and click Next.

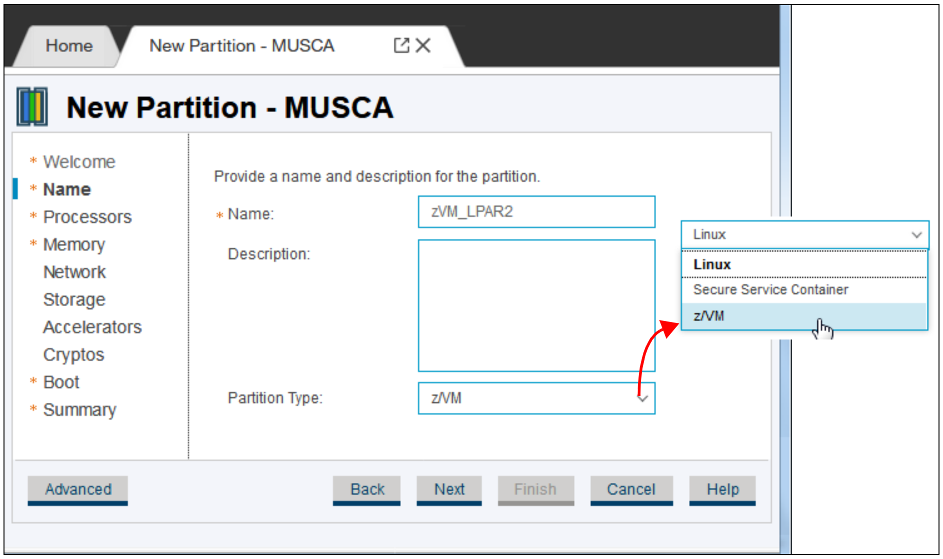

Figure 16-44 Partition name, description, and type

3. Make the following selections:
   - Processor resources (similar to Linux partition)
   - Memory resources (similar to Linux partition)
   - Network adapters and configure network devices (similar to Linux partition)
   - Host bus adapter and configure storage devices (similar to Linux partition)
4. Configure the accelerator adapter (zEDC) to our partition, as shown in Figure 16-45. Select the wanted adapter click **OK**.

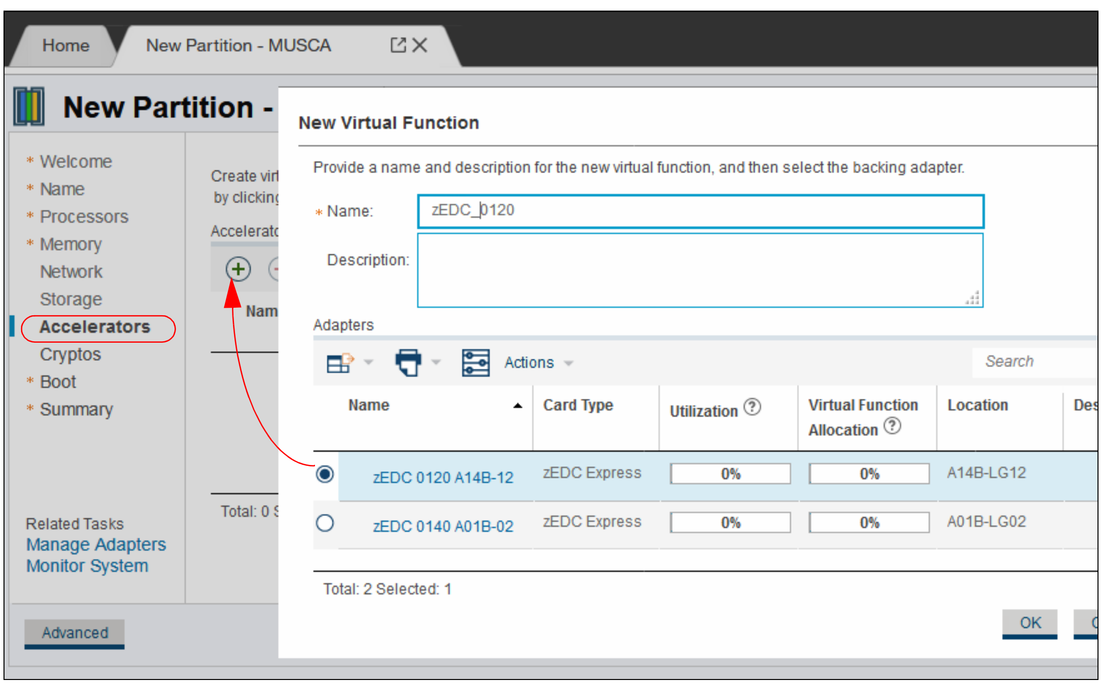

Figure 16-45 Selecting an accelerator adapter for the partition

5. An accelerator device (0001) is created, as shown in Figure 16-46. Click **Next** to configure a Crypto adapter.

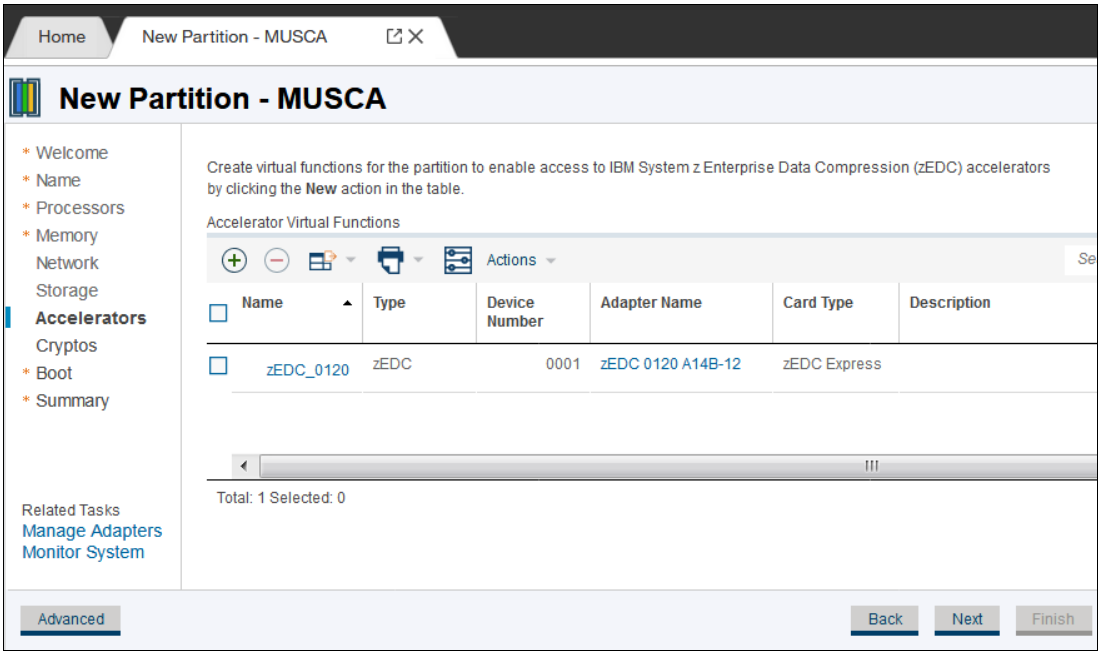

Figure 16-46 Accelerator device (data compression)

Figure 16-47 shows selecting the Crypto adapter.

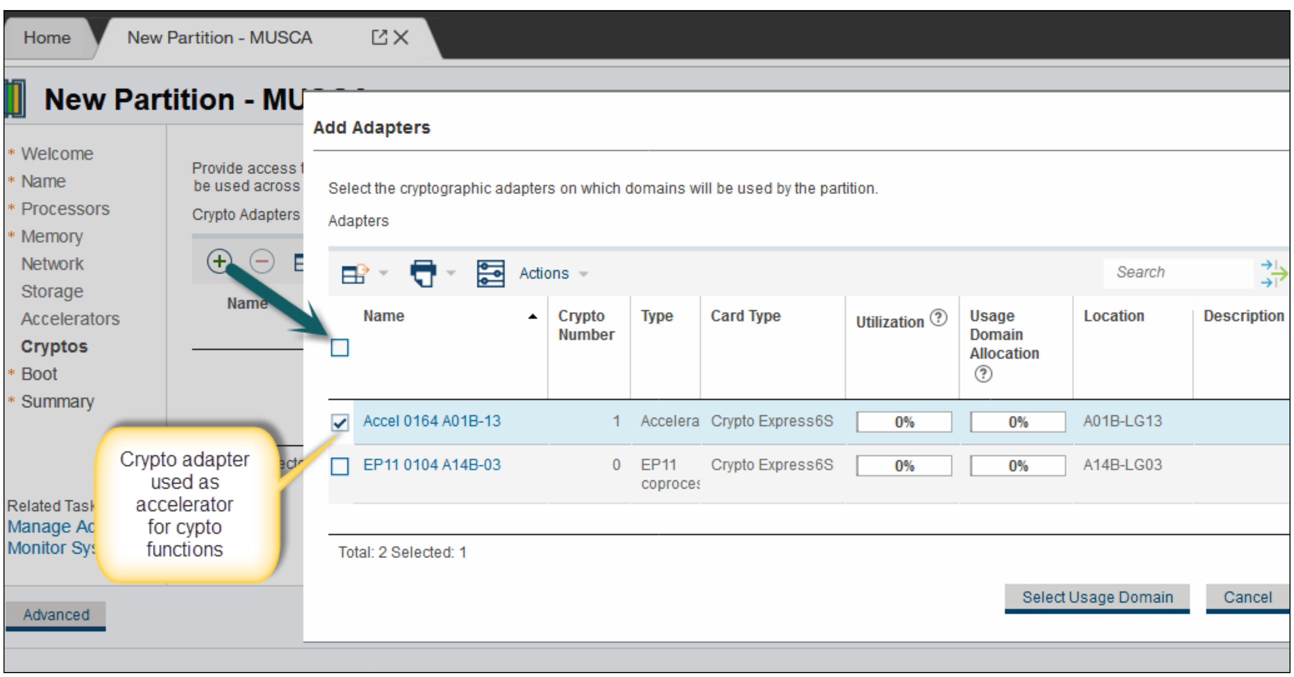

Figure 16-47 Cryptographic adapter selection (Crypto selected as an accelerator)

A cryptographic domain must be selected for the configured partition, as shown in Figure 16-48.

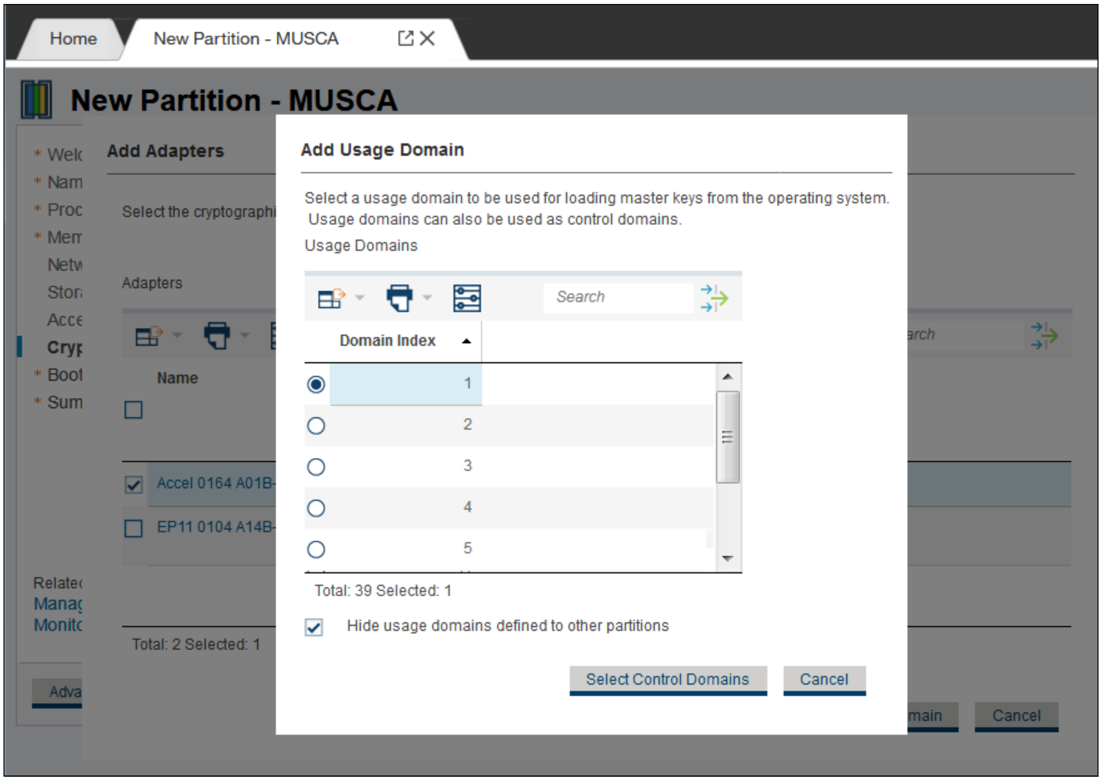

Figure 16-48 Selecting cryptographic domain for the selected adapter

6. Select the control domain, as shown in Figure 16-49.

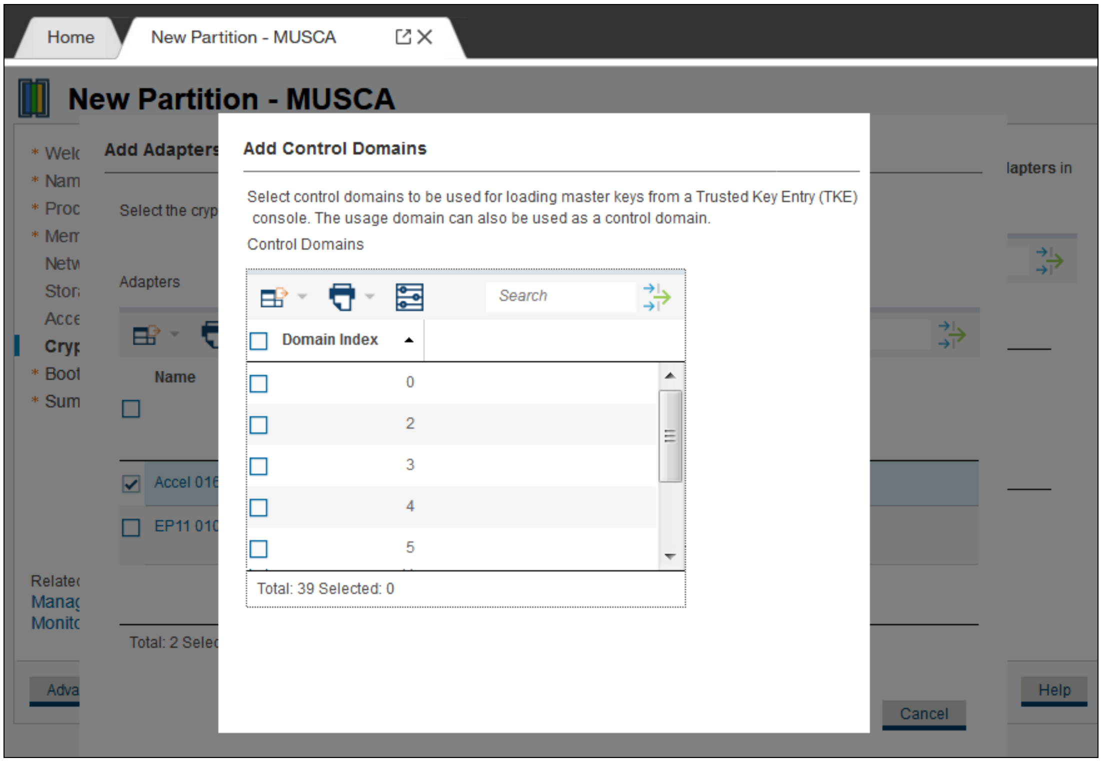

Figure 16-49 Control domain selection

Figure 16-50 shows the Crypto device configured for the z/VM partition.

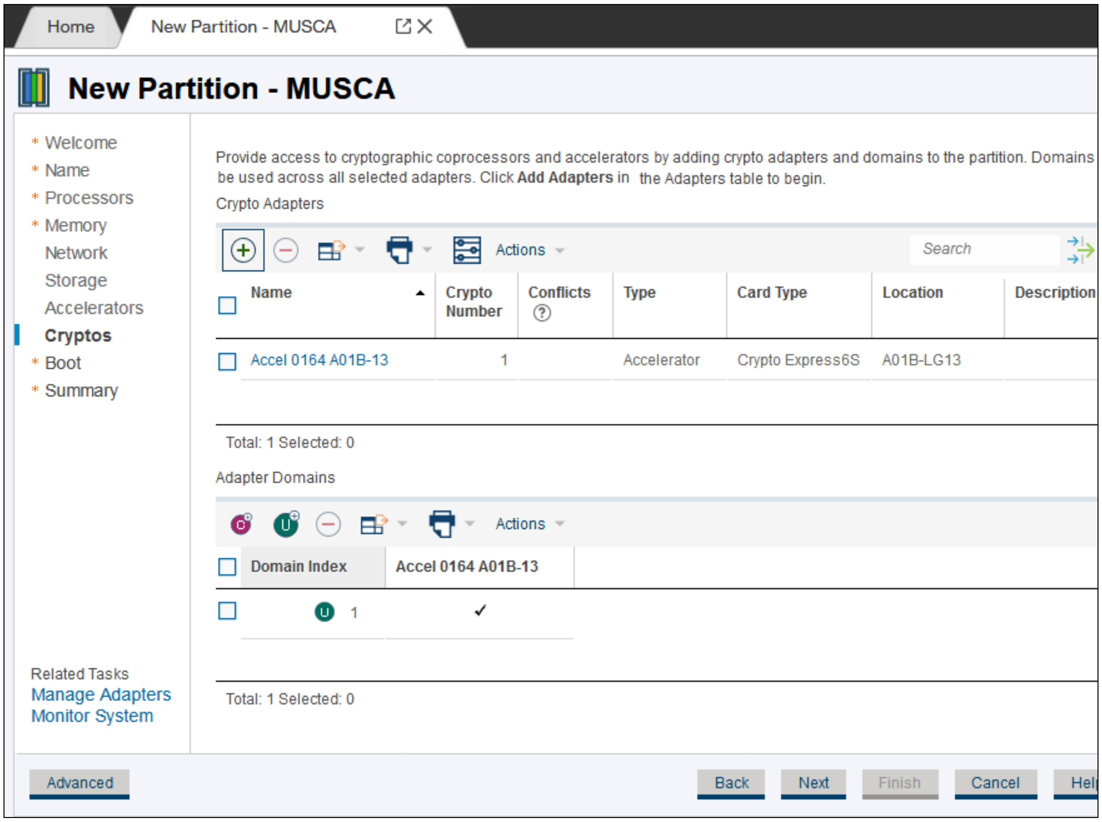

Figure 16-50 Crypto device configuration

7. Selecting the boot options (similar to what is shown in Figure 16-39 on page 413).
8. After selecting the partition boot options, finalize the partition creation by clicking **Finish** (see Figure 16-40 on page 413).
9. Save the information about the partition's storage device configuration by exporting the configured WWPNs. The WWPN information is required for SAN zoning and storage access (LUN Masking).

Note: This book does *not* describe how to install software in a partition. For more information about such installations, see your software of choice (operating system, hypervisor, or appliance) installation documentation.

16.3.5 Configuring a Secure Service Container partition

An SSC partition is configured similar to Linux or z/VM partitions. As such, we provide figures for the choices and features that are specific to this type of partition only.

Complete the following steps:

1. To configure a partition, start with the Systems Management interface on the HMC and select the CPC to be configured. Then, in the Tasks pane, select **New Partition**, as shown in Figure 16-42 on page 415.
2. Select the type of partition (**Secure Service Container**), as shown in Figure 16-51. For this type of partition, no operating system console is used, only web interface access. For more information about the user information configuration, see your documentation.

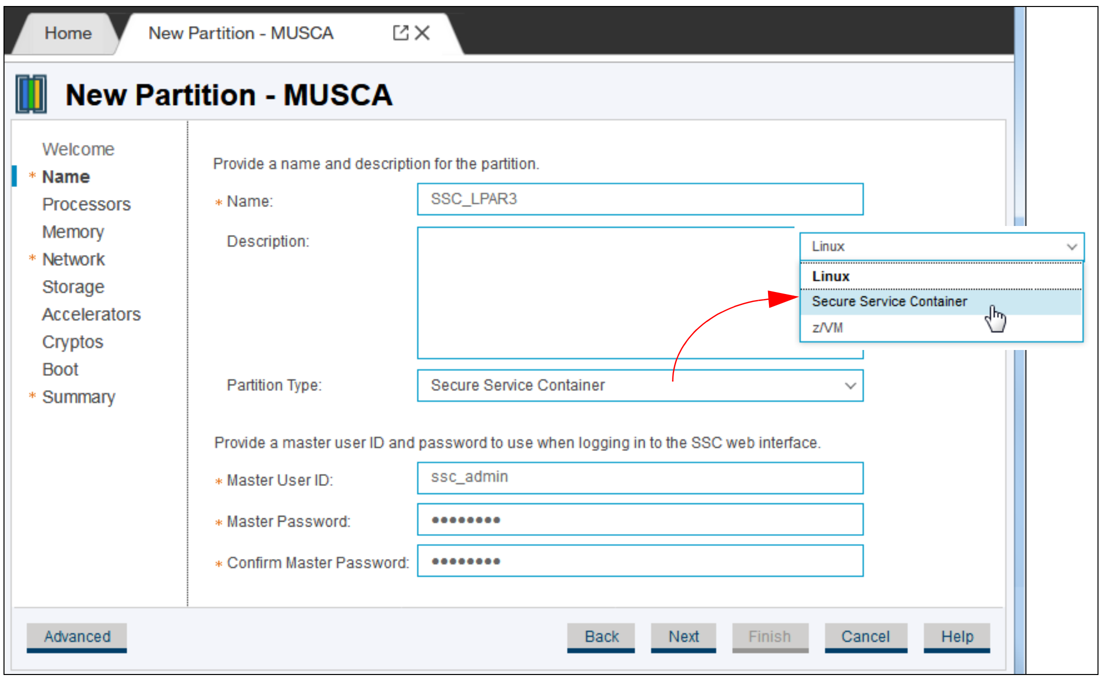

Figure 16-51 Secure Service Container partition description

3. Select the processor resources (CP or IFL).
4. Select the memory that is required for this partition (wanted and maximum).
5. Select the Network Interface Card adapter for this partition, as shown in Figure 16-52 on page 422.

> **Note:** A NIC device is required for this type of partition.

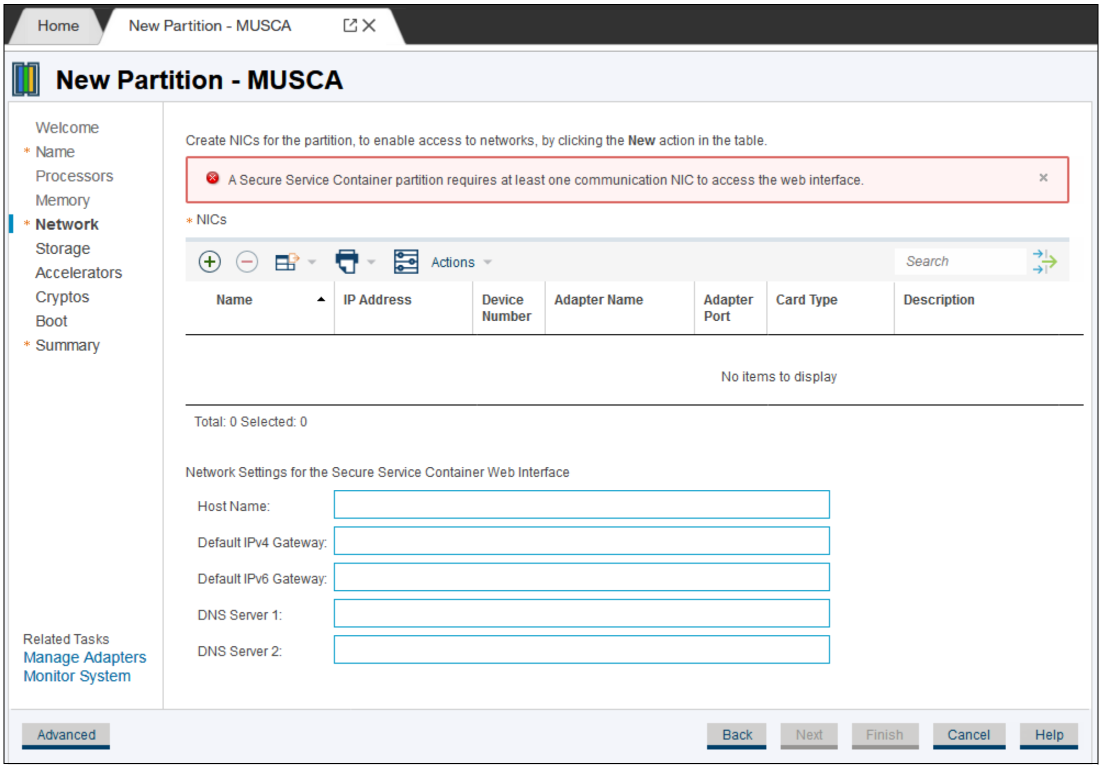

Figure 16-52 NIC selection

6. Configure TCP/IP information for the NIC (otherwise, the installer interface might not be accessible). The configuration we used is shown in Figure 16-53.

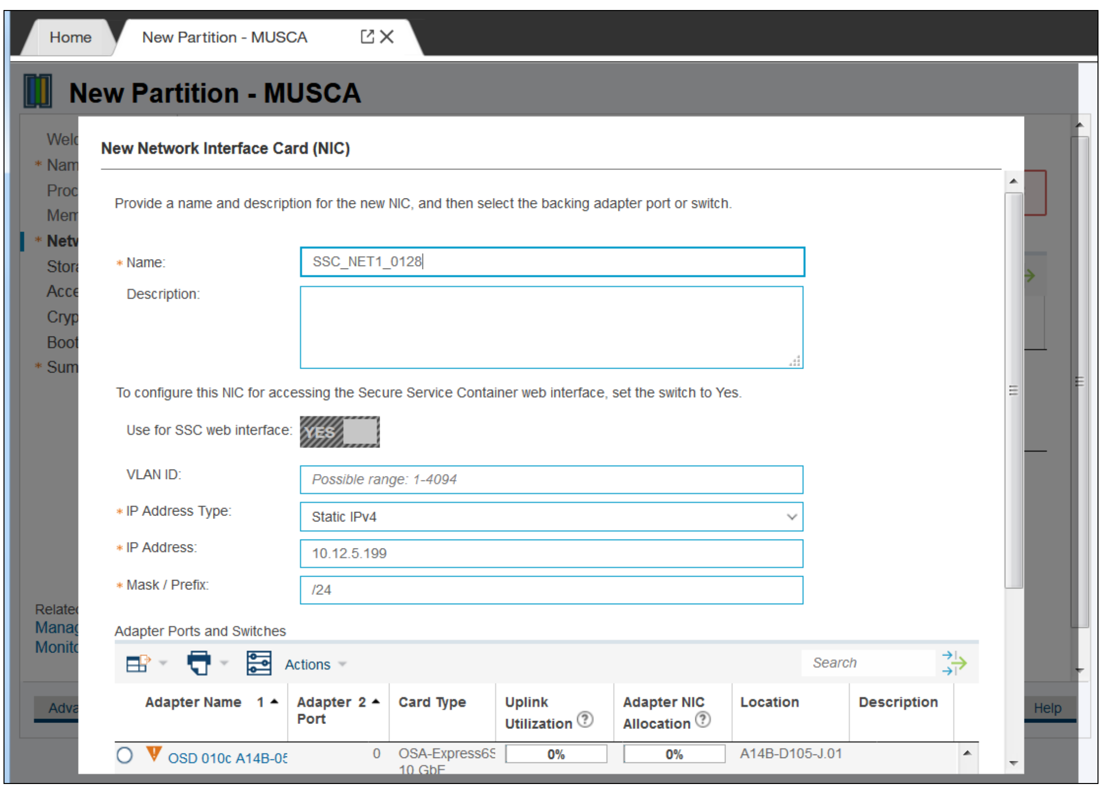

Figure 16-53 NIC description and IP information

7. Figure 16-54 shows the configured NIC. Click **Next** to configure the storage adapters.

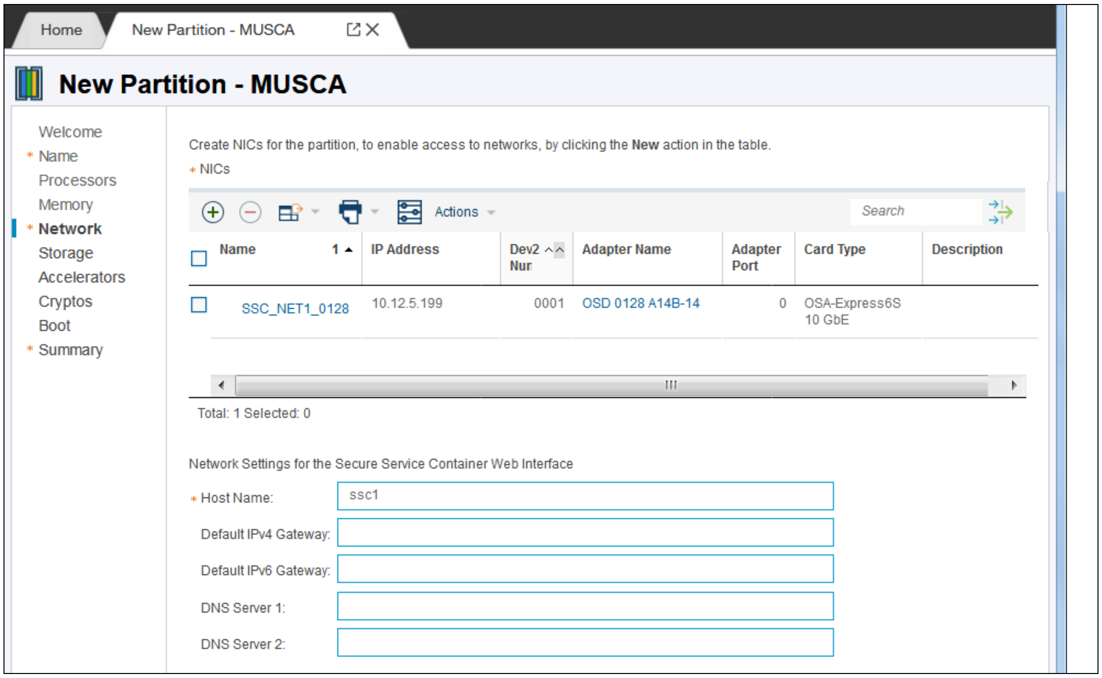

Figure 16-54 NIC configuration

8. Configure the storage adapter (HBA). In our example, we used FCP storage (see Figure 16-38 on page 412).
9. Optionally, you can configure Accelerators and Cryptos (if required by the code that runs in the partition).
10. After you see the available boot options, you can select to boot in installer mode (which is required if you want to install code on the storage that is allocated to the partition), as shown in Figure 16-55.

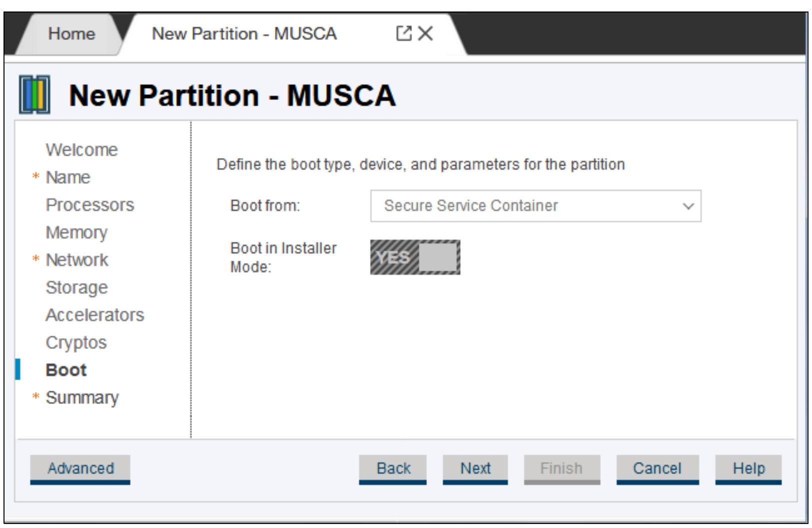

Figure 16-55 Boot options for SSC partition

11.The summary is shown in Figure 16-56. Click **Finish** to create the partition

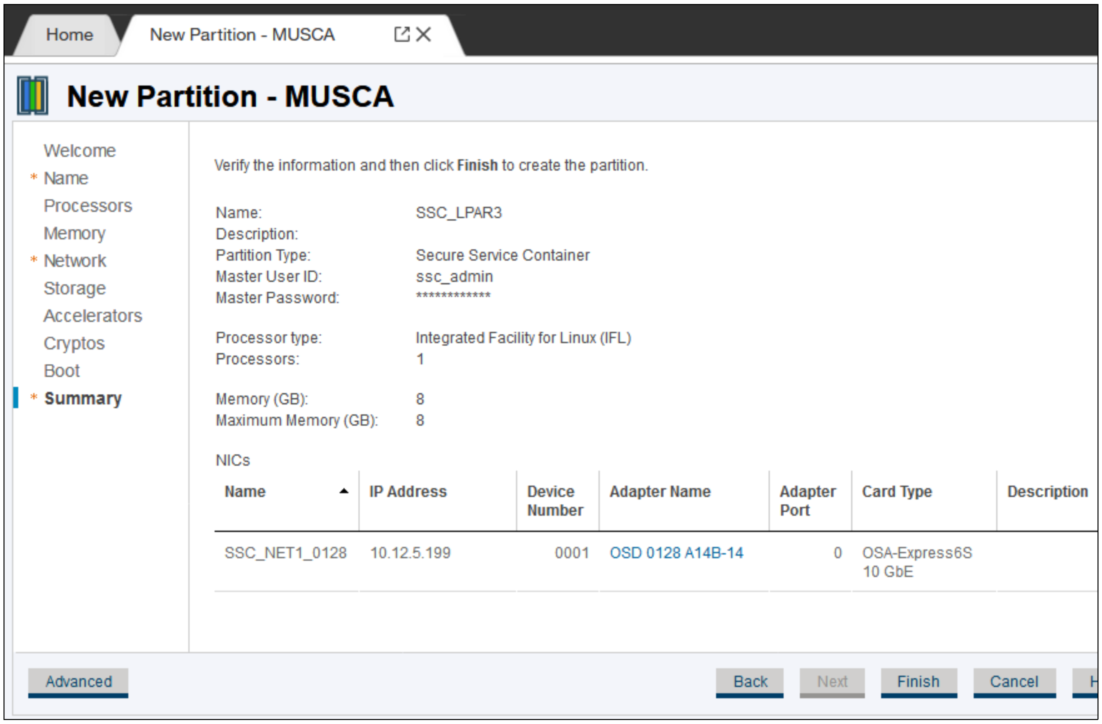

Figure 16-56 Secure Service Container partition summary

12.Save the information about the partition's storage device configuration by exporting the configured WWPNs. The WWPN information is required for SAN zoning and storage access (LUN Masking).

Note: This book does *not* describe how to install software in a partition. For more information about such installations, see your software of choice (operating system, hypervisor, or appliance) installation documentation.

16.3.6 Configuring partitions by using Advanced menus

In this scenario, we describe how to configure a Linux partition by using the Advanced menus. The Advanced menus allow a system administrator[5] who is familiar with the system to take a more granular approach in resource management and allocation.

For example, the Advanced menus allow configuring dedicated processor resources (only shared processor resources can be configured by using the standard (wizard) menus).

[5] System administrator role is equivalent to System Programmer (HMC role).

Complete the following steps:

1. Select the Advanced menus, as shown in Figure 16-57.

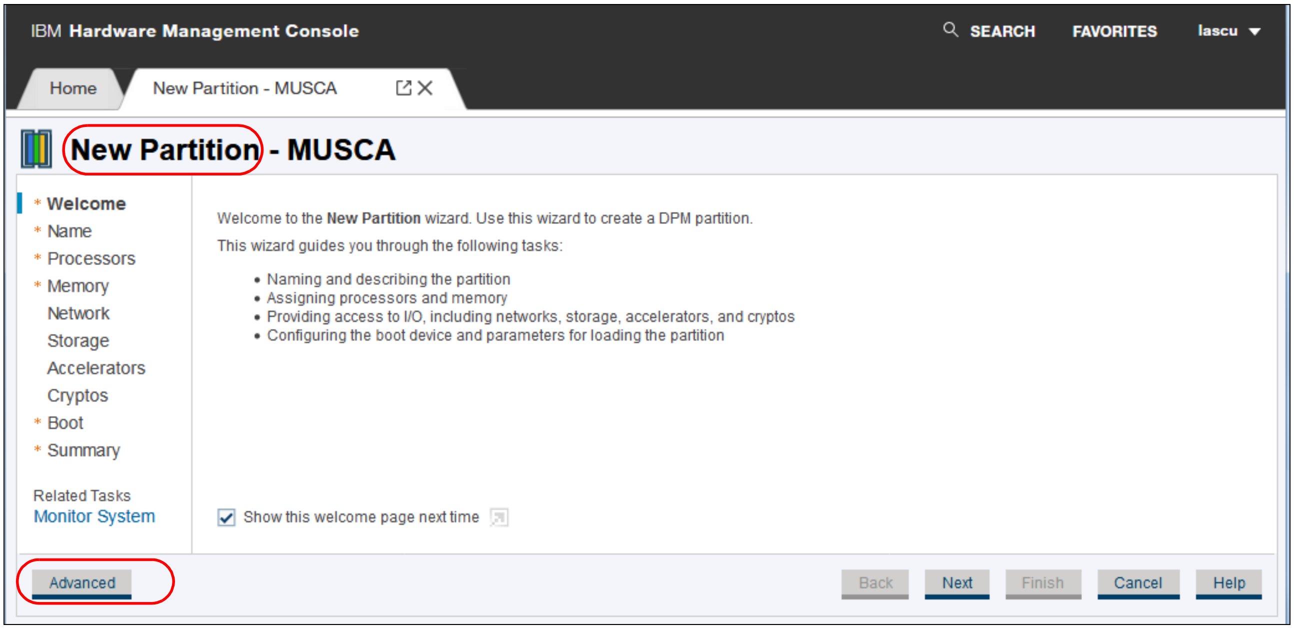

Figure 16-57 Initial menu for configuring partitions - selecting Advanced menus

2. Confirm by clicking **Switch**, as shown in Figure 16-58.

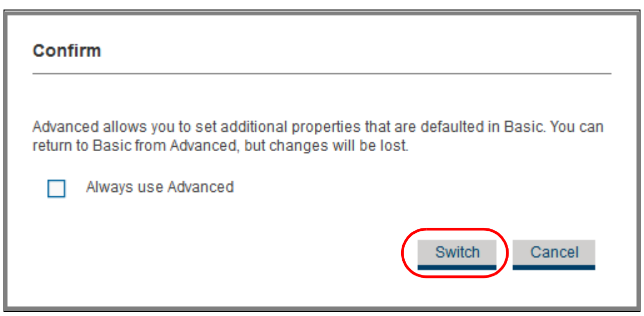

Figure 16-58 Advanced mode switch confirmation

3. In the Advanced menu (see Figure 16-59) window, scroll down to configure all resources.

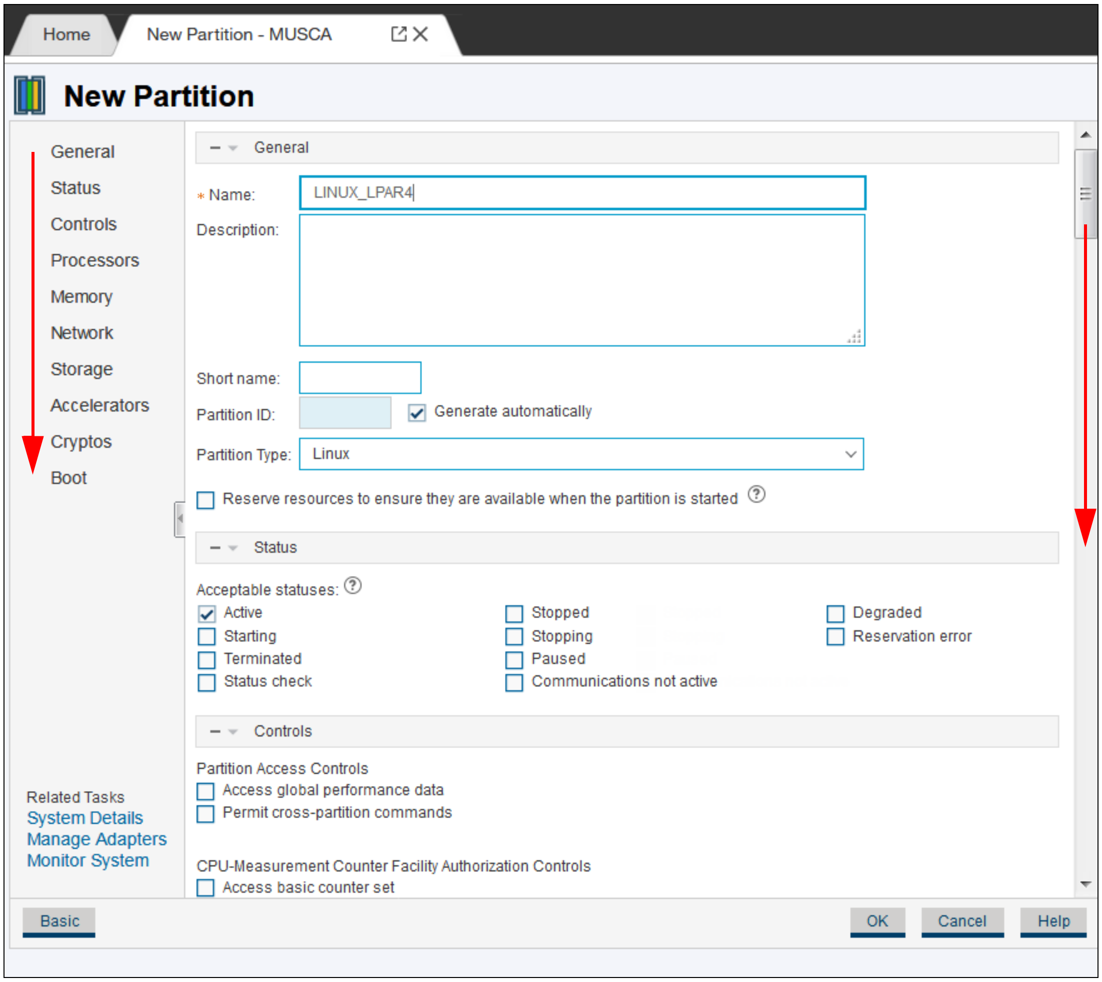

Figure 16-59 New Partition - Advanced menu

Because these configuration steps are similar to the standard and wizard mode, we do not cover these steps here.

16.3.7 Changing partition properties

After the partition is configured, its characteristics can be changed by using the Partition Details menu that is available in the Systems Details tab, as shown in Figure 16-60.

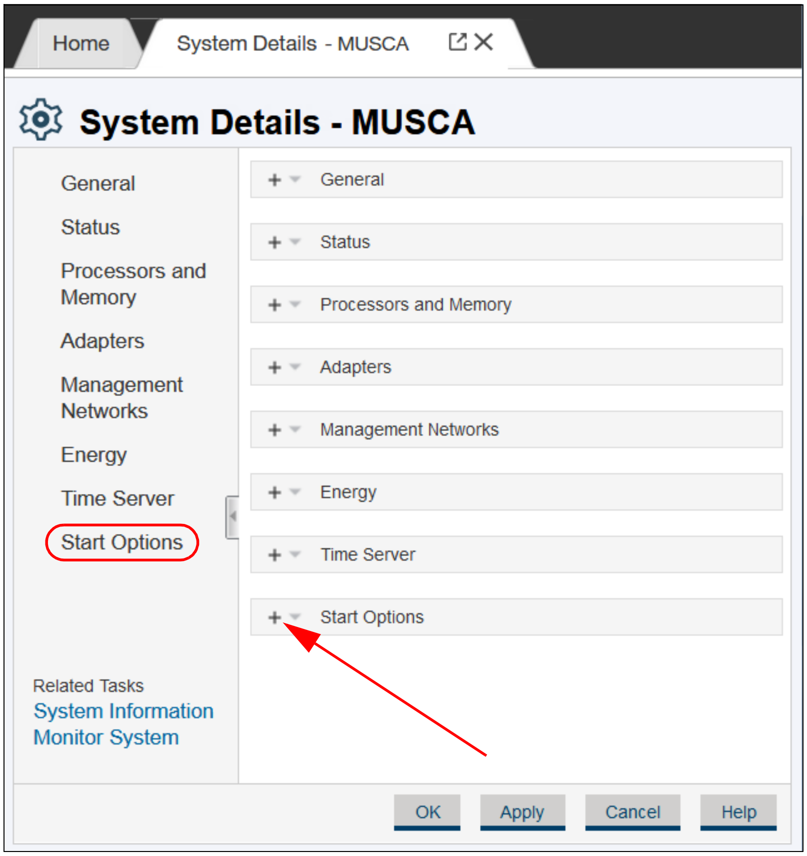

Figure 16-60 System Details tab

In the System Details tab, click **Start Options** (or scroll down to the Start Options section). Then, select the wanted partition and open the Partition Details tab.

The extra configuration parameters that are not presented when a new partition is configured by using the standard and wizard configuration menu can be changed after the partition is created.

The Partition Details tab (see Figure 16-61) is used to browse the partition configuration.

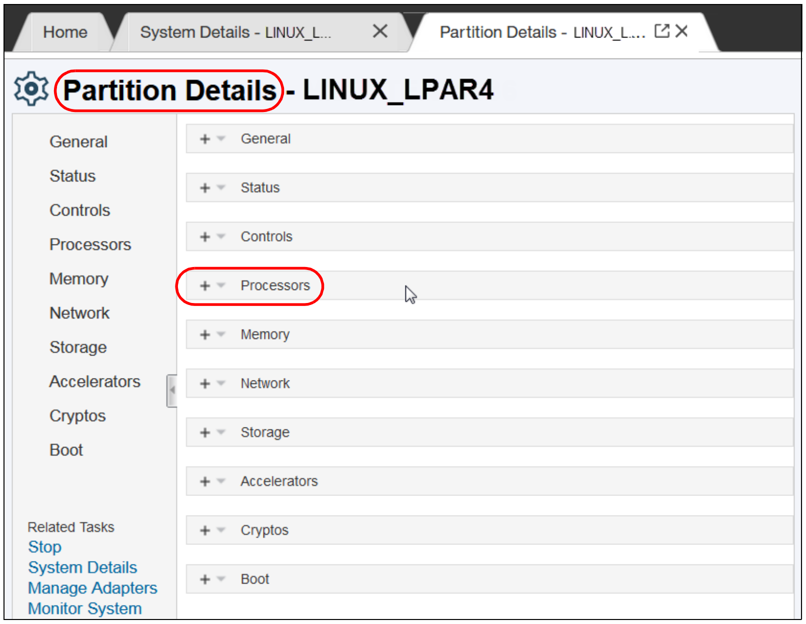

Figure 16-61 Partition details

Some of the other partition resources and properties that can be configured by using the Advanced menu include *Dedicated* processors, as shown in Figure 16-62.

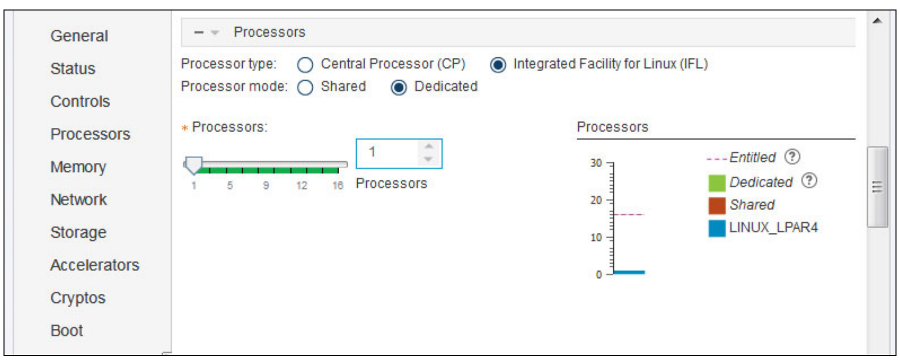

Figure 16-62 Configuring Dedicated processors for a partition

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 material

The web material that is associated with this book is available in softcopy on the internet from the IBM Redbooks web server:

ftp://www.redbooks.ibm.com/redbooks/SG248560

Alternatively, you can visit the IBM Redbooks website:

ibm.com/redbooks

Search for SG248560, select the title, and then, click **Additional materials** to open the directory that corresponds with the IBM Redbooks form number, SG248560.

Using the web material

Additional web material that accompanies this book includes the following file:

File name	*Description*
8560_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 `.zip` file into this folder.

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications that are referenced in this list might be available in softcopy only:

- *FICON CTC Implementation*, REDP-0158
- *FICON Planning and Implementation Guide*, SG24-6497
- *IBM Z Connectivity Handbook*, SG24-5444
- *IBM z/OS V2R2 Communications Server TCP/IP Implementation Volume 1: Base Functions, Connectivity, and Routing*, SG24-8360
- *IBM z13 Configuration Setup*, SG24-8260
- *IBM z14 Technical Guide, SG24-8451*
- *IBM z14 Technical Introduction*, SG24-8450
- *I/O Configuration Using z/OS HCD and HCM*, SG24-7804
- *Mainframe from Scratch: Hardware Configuration and z/OS Build*, SG24-8329
- *OSA-Express Implementation Guide*, SG24-5948
- *OSA-Express Integrated Console Controller Implementation Guide*, SG24-6364
- *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259
- *Server Time Protocol Implementation Guide*, SG24-7281
- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Recovery Guide*, SG24-7380

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

The following websites are also relevant as further information sources:

- Capping Technologies and 4HRA Optimization document:

  ftp://public.dhe.ibm.com/eserver/zseries/zos/wlm/Capping_Technologies_and_4HRA_Optimization_2016.pdf
- CFSizer tool:

  https://www.ibm.com/systems/support/z/cfsizer/

- CTC Definitions for z Systems:
  http://www.vm.ibm.com/devpages/altmarka/ctc.html
- Hardware Configuration (HCD and HCM):
  https://www.ibm.com/systems/z/os/zos/features/hcm/
- IBM Knowledge Center:
  https://www.ibm.com/support/knowledgecenter
- IBM Resource Link:
  https://www.ibm.com/servers/resourcelink/svc03100.nsf?OpenDatabase
- IBM z Systems Batch Network Analyzer (zBNA) Tool website:
  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132
- IBM z/OS SMC Applicability Tool (SMC-AT) document:
  ftp://public.dhe.ibm.com/software/os/systemz/pdf/SMC_Applicability_Tool_Overview_03-10-16.pdf
- IBM's Shared Memory Communications over RDMA (SMC-R) Protocol, RFC 7609:
  https://tools.ietf.org/html/rfc7609
- Parallel Sysplex white paper page:
  https://www.ibm.com/systems/z/advantages/pso/whitepaper.html
- Shared Memory Communications Reference Information website:
  https://www.ibm.com/software/network/commserver/SMC/
- Techdocs website:
  https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3659
- WWPN Prediction Tool:
  https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/wwpnMain?OpenDocument&pathID=
- z/OS downloads page:
  https://www.ibm.com/systems/z/os/zos/tools/downloads/index.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Redbooks

IBM z14 Model ZR1 Configuration Setup

SG24-8560-00

ISBN 073845821x

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

SG24-8560-00

ISBN 073845821x

Printed in U.S.A.