Redbooks
ibm.com/redbooks

IBM z13 Configuration Setup

Paolo Bruni

Tom Carielli

Mark Challen

Klaus Horn

Peter Hoyle

Kazuhiro Nakajima

Martin Soellig

z Systems

IBM

Redbooks

International Technical Support Organization

IBM z13 Configuration Setup

April 2015

SG24-8260-00

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

First Edition (April 2015)

This edition applies to the implementation of the IBM z13 (machine type 2964)

Contents

Notices

This information was developed for products and services offered in the U.S.A.

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to non-IBM 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 supply in any way it believes appropriate without incurring any obligation to you.

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

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

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

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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DataPower®
DB2®
DS8000®
Electronic Service Agent™
FICON®
FlashCopy®
GDPS®
HiperSockets™
HyperSwap®
IBM®
IBM Resiliency Services™
IBM z Systems™
IBM z13™
IMS™
Language Environment®
MQSeries®
MVS™
Parallel Sysplex®
POWER7®
PR/SM™
Processor Resource/Systems Manager™
Redbooks®
Redbooks (logo) ®
Resource Link®
RETAIN®
RMF™
System p®
System Storage®
System z®
System z10®
System z9®
VTAM®
WebSphere®
z Systems™
z/Architecture®
z/OS®
z/VM®
z/VSE®
z10™
z13™
z9®
zEnterprise®

The following terms are trademarks of other companies:

Inc., and Inc. device are trademarks or registered trademarks of Kenexa, an IBM Company.

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

This IBM® Redbooks® publication helps you install, configure, and maintain the IBM z13™. The z13 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 Systems™ configuration and implementation. Readers should be generally familiar with current IBM z Systems technology and terminology. For details about the functions of the z13, see *IBM z13 Technical Introduction*, SG24-8250 and *IBM z13 Technical Guide*, SG24-8251.

Authors

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

Paolo Bruni is an ITSO Project Leader and is based at the Silicon Valley Lab in San Jose, CA. Since 1998, Paolo has authored Redbooks publications about IBM IMS™, DB2® for z/OS®, and related tools, and has conducted workshops worldwide. During his many years with IBM in development and in the field, Paolo's work has been related mostly to IBM z Systems.

Tom Carielli is a veteran Systems Programmer in Poughkeepsie, NY, who has over 35 years of experience with IBM mainframes as a Systems and Application Programmer. Tom has a Master's Degree in Computer Science from the University of Michigan. He works closely with customers from all around the world to help support z Systems hardware and software. Tom has been part of the z Systems Client Center team in Poughkeepsie since its inception in 1997. His responsibilities include the hardware configuration and operating systems setup of customer environments to benchmark, tune, and test applications, measure performance, and determine workload capacity.

Mark Challen works for IBM Resiliency Services™ in the United Kingdom. He has over 26 years of experience working for IBM in the mainframe field, the past 14 of those as a z/OS Systems Programmer. His areas of expertise include SMP/e, HCD AND REXX.

Klaus Horn is an IT Specialist in the TSCC Hardware FE z Systems center in Mainz, Germany. He has worked at IBM for 14 years and has a total of 8 years of experience in supporting mainframe clients. His areas of expertise include z Systems installation and installation planning.

Peter Hoyle works for IBM as a z Systems configuration and solution designer in Australia. He has more than 35 years of experience in the IT industry, including 15 years in IBM mainframe configuration management and design. He has worked at IBM since 1999. His areas of expertise include HCM and HCD, and implementing configuration solutions, creating diagrams and documentation. He was a co-author of several z Systems I/O configuration Redbooks publications from the IBM System z9 to the IBM z13.

Kazuhiro Nakajima is a Senior IT Specialist in IBM Japan. He has over 10 years of experience in technical sales and supporting z Systems clients. His areas of expertise include z Systems hardware, z/OS, z Systems connectivity. He has also co-authored the IBM zEnterprise EC12 Configuration Setup.

Martin Soellig is a Consultant IT Specialist in Germany. He has 25 years of experience working in the z Systems field. He holds a degree in Mathematics from University of Hamburg. His areas of expertise include z/OS and z Systems hardware, specifically in IBM Parallel Sysplex® and GDPS® environments.

The authors in Poughkeepsie. From left to right: Kazuhiro, Paolo, Tom, Martin, Mark, Klaus, and Peter (photo courtesy of Rich Conway)

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

Octavian Lascu
Klaus Meissner
Bill White

Thanks to the authors of the *IBM zEnterprise EC12 Configuration Setup*, SG24-8034, published in May 2013:

Sandra Tomoko Minami Abe
Srikanth R. Adige
Lei Ming Cao
Joao Francisco Veiga Kiffer
Woojin Lee
Gianmauro De Marchi
Klaus Meissner
Kazuhiro Nakajima
Karan Singh
Botrous Thalouth

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:

  https://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

Introducing the IBM z13

The IBM z13, like its predecessors, is designed from the chip level up to support data processing. This includes a strong, fast I/O infrastructure, cache on the chip to bring data close to processing power, security and compression capabilities of the coprocessors and I/O features, and the 99.999% data availability design of the coupling technologies.

In this chapter, we provide a quick overview of the new member of the IBM z Systems family. For more specific details, see *IBM z13 Technical Guide*, SG24-8251.

This chapter includes the following topics:

- IBM z13 technical overview
- Hardware Management Console and Support Element
- IBM z BladeCenter Extension (zBX) Model 004

1.1 IBM z13 technical overview

This section briefly reviews the major elements of the z13:

- Models
- Model upgrade paths
- Frames
- CPC drawer
- I/O connectivity: PCIe and InfiniBand
- I/O subsystems
- Coupling Links support
- Special purpose features:
  - Cryptography
  - Flash Express
  - zEDC Express

1.1.1 Models

The z13 has a machine type of 2964. Five models are offered: N30, N63, N96, NC9 and NE1. The model name indicates the maximum number of processor units (PUs) available for purchase ("C9" stands for 129 and "E1" for 141). A PU is the generic term for the IBM z/Architecture® processor unit.

With the z13, some PUs are part of the system base, that is, they are not part of the PUs that can be purchased by clients. They are characterized by default:

- System assist processor (SAP) used by the channel subsystem. The number of predefined SAPs is dependent on the z13 model.
- One integrated firmware processor (IFP). The IFP is used in support of select features, such as zEDC and 10 Gigabit Ethernet (10GbE) Remote Direct Memory Access (RDMA) over Converged Ethernet (10GbE RoCE).
- Two spare PUs that can transparently assume any characterization in a permanent failure of another PU.

The PUs that clients can purchase can assume any of the following characterizations:

- Central processor (CP) for general-purpose use.
- Integrated Facility for Linux (IFL) for the exploitation of Linux on z Systems.
- z Systems Integrated Information Processor (zIIP). One CP must be installed with or before the installation of any zIIPs.

> **zIIPs:** At least one CP must be purchased with, or before, a zIIP can be purchased. Clients can purchase up to two zIIPs for each purchased CP (assigned or unassigned) on the system.

- Internal Coupling Facility (ICF) to be used by the Coupling Facility Control Code (CFCC).
- Additional system assist processor (SAP) to be used by the channel subsystem.

> **Support:** z13 does not support zAAPs.

1.1.2 Model upgrade paths

Any z13 can be upgraded to another z13 hardware model. All upgrades from Models N30, N63, and N96, to Models N63, N96, and NC9 respectively are concurrent. The upgrades to NE1 are disruptive (that is, the system is unavailable during the upgrade). Any zEC12 or z196 model can be upgraded to any z13 model, which is also disruptive. See Chapter 5, "Upgrading from an IBM zEC12 to IBM z13" on page 159

> **Note:** An air-cooled z13 cannot be converted to a water-cooled z13, and vice versa.

The following items are not supported:

- Downgrades within the z13 models
- Upgrade from a zBC12 or z114 to z13
- Upgrades from IBM System z10 or earlier platforms
- Attachment of an IBM zEnterprise BladeCenter Extension (zBX) Model 002 or Model 003 to a z13

1.1.3 Frames

The z13 has two frames, which are bolted together and are known as the A frame and the Z frame. The frames contain the following CPC components:

- Up to four CPC drawers located in frame A.
- PCIe I/O drawers, and I/O drawers, which hold I/O features and special purpose features
- Power supplies
- An optional Internal Battery Feature (IBF)
- Cooling units for either air or water cooling
- Two new System Control Hubs to interconnect the CPC components via Ethernet
- Two new 1U rack-mounted Support Elements with their keyboards, mouses, and displays mounted on a tray in the z frame.

1.1.4 CPC drawer

Up to four CPC drawers are installed in frame A of the z13. Each CPC drawer houses the SCMs, memory, and I/O interconnects.

SCM technology

The z13 is built on the proven superscalar microprocessor architecture of its predecessor, and provides several enhancements over the zEC12. Each z13 is physically divided in two nodes. Each node has four single chip modules (SCMs), three processor unit (PU) SCMs and one storage control (SC) SCM, therefore the CPC drawer has six PU SCMs and two SC SCMs. The PU SCM has eight core, with six, seven, or eight active cores, which can be characterized as CPs, IFLs, ICFs, zIIPs, or system assist processors (SAPs). Two CPC drawers sizes are offered: 39 and 42 cores.

The SCM provides a significant increase in system scalability and an additional opportunity for server consolidation. All CPC drawers are interconnected for two high-speed communication links cable-based, in a full star topology, through the L4 cache. This configuration allows the system to be operated and controlled by the IBM Processor Resource/Systems Manager™ (PR/SM™) facility as a memory-coherent and cache-coherent SMP.

The PU configuration includes two designated spare PUs per CPC and a variable number of SAPs. The SAPs scale with the number of CPC drawers that are installed in the server. For example, there are six standard SAPs with one z13 installed, and up to 24 when four z13 are installed. In addition, one PU is used as an integrated firmware processor (IFP) and is not available for client use. The remaining PUs can be characterized as CPs, IFL processors, zIIPs, ICF processors, or additional SAPs.

Processor features

The processor chip has an eight-core design, with either six, seven, or eight active cores, and operates at 5 GHz. Each core on the PU chip includes an enhanced dedicated coprocessor for data compression and cryptographic functions, such as the Central Processor Assist for Cryptographic Function (CPACF). The cryptographic performance of CPACF has improved up to two times over the zEC12.

Having standard clear key cryptographic coprocessors that are integrated with the processor provides high-speed cryptography for protecting data.

Hardware data compression can play a significant role in improving performance and saving costs over performing compression in software. The zEDC Express feature offers additional performance and savings over the coprocessor. Their functions are not interchangeable.

The micro-architecture of core has altered radically to increase parallelism and improve the pipeline. The core has a new branch prediction and instruction fetch front end to support simultaneous multithreading in a single core and to improve the branch prediction throughput, a wider instruction decode (six instruction per cycle) and ten arithmetic logical execution units, that offer double instruction bandwidth over the zEC12.

Each core has two hardware decimal floating point unit that is designed according to a standardized, open algorithm. Much of today's commercial computing is decimal floating point, so two on-core hardware decimal floating point meets the requirements of business and user applications, provides greater floating point execution and improved performance, precision, and function.

In the unlikely case of a permanent core failure, each core can be individually replaced by one of the available spares. Core sparing is transparent to the operating system and applications.

Simultaneous multithreading (SMT)

The micro-architecture of the core of the z13 permits the execution of simultaneous multithreading (SMT) in the same zIIP or IFL core. Two threads can run in parallel in a single core, dynamically sharing processor resources such as execution units and caches. This facility permits a more efficient utilization of the core and increased capacity, because while one of the threads is waiting for a storage access (cache miss), the other thread that is executing simultaneously in the core can utilize the shared resources rather than remain idle.

Single-instruction, multiple-data (SIMD) instruction set

The z13 instruction set architecture includes a subset of 139 new instructions, such as *single-instruction, multiple-data (SIMD)*, were added to improve efficiency of complex mathematical models and vector processing. These new instructions allow a lager number of operands to be processed with a single instruction. The SIMD instructions utilize the superscalar core to process operands in parallel enabling more interactions.

Transactional Execution facility

The z13 like the previous machine has a set of instructions that allows defining groups of instructions that are run atomically. That is, either all the results are committed or none are.

The facility provides for faster and more scalable multithreaded execution, and is known as *hardware transactional memory*.

Out-of-order (OOO) execution

As in its predecessor zEC12, z13 has an enhanced superscalar microprocessor with out-of-order (OOO) execution to achieve faster throughput. With OOO, instructions might not run in the original program order, although results are presented in the original order. For instance, OOO allows a few instructions to complete while another instruction is waiting. Up to six instructions can be decoded per system cycle, and up to ten instructions can be in execution.

Concurrent processor unit (PU) conversions

The z13 supports concurrent conversion between various PU types, providing flexibility to meet changing business environments. CPs, IFLs, zIIPs, ICFs, and optional SAPs can be converted to CPs, IFLs, zIIPs, ICFs, and optional SAPs.

Memory subsystem and topology

z13 uses a new buffered dual inline memory modules (DIMM) technology. For this purpose, IBM has developed a chip that controls communication with the PU, and drives address and control from DIMM to DIMM. The DIMM capacities are 16 GB, 32 GB, 64 GB and 128 GB.

Memory topology provides the following benefits:

- Redundant array of independent memory (RAIM) for protection at the dynamic random access memory (DRAM), DIMM, and memory channel levels
- A maximum of 10 TB of user configurable memory with a maximum of 12.5 TB of physical memory (with a maximum of 10 TB configurable to a single LPAR)
- One memory port for each PU chip, and up to five independent memory ports per z13
- Increased bandwidth between memory and I/O
- Asymmetrical memory size and DRAM technology across CPC drawer
- Large memory pages (1 MB and 2 GB)
- Key storage
- Storage protection key array that is kept in physical memory
- Storage protection (memory) key is also kept in every L2 and L3 cache directory entry
- A larger (96 GB) fixed-size HSA that eliminates having to plan for HSA

PCIe fanout hot-plug

The *PCIe fanout* provides the path for data between memory and the PCIe I/O cards through the PCIe 16 GBps bus and cables. The PCIe fanout is hot-pluggable. In an outage, a redundant I/O interconnect allows a PCIe fanout to be concurrently repaired without loss of access to its associated I/O domains. Up to ten PCIe fanouts are available per CPC processor drawer. The PCIe fanout can also be used for the Integrated Coupling Adapter (ICA). As long as redundancy in coupling link connectivity is ensured, the PCIe fanout can be concurrently repaired.

Host channel adapter (HCA) fanout hot-plug

The host channel adapter (HCA) fanout provides the path for data between memory and the I/O cards residing in an I/O drawer through 6 GBps InfiniBand (IFB) cables. The HCA fanout is hot-pluggable. In an outage, an HCA fanout can be concurrently repaired without the loss of access to its associated I/O features, using redundant I/O interconnect. InfiniBand optical

HCA3-O and HCA3-O LR, used to provide connectivity between members of a sysplex are orderable features in z13, and can also be carried forward on a miscellaneous equipment specification (MES) from a zEC12 or z196. Up to four HCA fanouts are available per CPC drawer. The HCA fanout can also be used for the InfiniBand coupling links (HCA3-O and HCA3-O LR). As long as redundancy in coupling link connectivity is ensured, the HCA fanout can be concurrently repaired.

1.1.5 I/O connectivity: PCIe and InfiniBand

The z13 offers various improved features and exploits technologies, such as PCIe, InfiniBand, and Ethernet. This section briefly reviews the most relevant I/O capabilities.

The z13 takes advantage of PCIe Generation 3 to implement the following features:

- PCIe Gen3 (Generation 3) fanouts that provide 16 GBps connections to the PCIe I/O features.
- PCIe Gen3 (Generation 3) fanouts that provide 8 GBps coupling link connections through the new IBM Integrated Coupling Adapter (ICA) or PCIe-O SR.

The z13 takes advantage of InfiniBand to implement the following features:

- A 6 GBps I/O bus that includes the InfiniBand infrastructure (HCA2-C) for the I/O drawer for non-PCIe features. Note, HCA2-C is only available as a carry forward.
- Parallel Sysplex coupling links using IFB: 12x InfiniBand coupling links (HCA3-O) for local connections and 1x InfiniBand coupling links (HCA3-O LR) for extended distance connections between any two zEnterprise CPCs. The 12x IFB link (HCA3-O) has a bandwidth of 6 GBps and the 1X InfiniBand links have a bandwidth of 5 Gbps.

1.1.6 I/O subsystems

The z13 I/O subsystem is similar to the one on zEC12 and includes a new PCIe Gen-3 infrastructure. The I/O subsystem is supported by both a PCIe bus and an I/O bus similar to that of zEC12. It includes the InfiniBand Double Data Rate (IB-DDR) infrastructure, which replaced the self-timed interconnect that was used in previous z Systems platforms. This infrastructure is designed to reduce processor usage and latency, and provide increased throughput.

z13 offers two I/O infrastructure elements for holding the I/O features: PCIe I/O drawers, for PCIe features; and up to two I/O drawers for non-PCIe features.

PCIe I/O drawer

The *PCIe I/O drawer*, together with the PCIe I/O features, offers improved granularity and capacity over previous I/O infrastructures. It can be concurrently added and removed in the field, easing planning. Only PCIe cards (features) are supported, in any combination. Up to five PCIe I/O drawers can be installed on a z13.

I/O drawer

On the z13, I/O drawers are supported only when carried forward on upgrades from zEC12 or z196 to z13. For a new z13 order, it is not possible to order an I/O drawer.

The z13 can have up to two I/O drawers. I/O drawers can accommodate up to eight FICON Express8 features only. Based on the number of I/O features that are carried forward, the configuration determines the number of required I/O drawers.

Native PCIe and integrated firmware processor (IFP)

Native PCIe was introduced with the zEDC and RoCE Express features, which are managed differently from the traditional PCIe features. The device drivers for these adapters are available in the operating system. The diagnostics for the adapter layer functions of the native PCIe features are managed by LIC that is designated as a resource group partition, which runs on the integrated firmware processor (IFP). For availability, two resource groups are present and share the IFP.

During the ordering process of the native PCIe features, features of the same type are evenly spread across the two resource groups (RG1 and RG2) for availability and serviceability reasons. Resource groups are automatically activated when these features are present in the CPC.

I/O and special purpose features

The z13 supports the following PCIe features on a new build, which can only be installed in the PCIe I/O drawers:

- FICON Express16S Short Wave (SX) and 10 KM Long Wave (LX) (6.2 miles)
- FICON Express 8S Short Wave (SX) and 10 KM Long Wave (LX) (6.2 miles)
- OSA-Express5S 10 GbE Long Reach (LR) and Short Reach (SR)
- OSA-Express5S, GbE LX and SX, and 1000BASE-T
- 10GbE RoCE Express
- Crypto Express5S
- Flash Express
- zEDC Express

When carried forward on an upgrade, the z13 also supports the following features in the PCIe I/O drawers:

- FICON Express8S 10 KM LX and SX
- OSA-Express5S
- 10GbE RoCE Express
- Flash Express
- zEDC Express

When carried forward on an upgrade, the z13 also supports up to two I/O drawers on which the following feature can be installed:

- FICON Express8 10 KM LX and SX

FICON channels

Up to 160 features with up to 320 FICON Express16S or 160 features up to 320 FICON Express 8S channels are supported. The FICON Express8S features support a link data rate of 2 Gbps, 4 Gbps, or 8 Gbps and the FICON Express 16S supports 4Gbs, 8Gbs, and 16Gbs.

The z13 FICON features support the following protocols:

- FICON (FC) and High Performance FICON for z Systems (zHPF). zHPF offers improved access to data, which is of special importance to OLTP applications.
- FICON Channel-to-channel (CTC).
- Fibre Channel Protocol (FCP).

FICON also offers the following capabilities:

- Modified Indirect Data Address Word (MIDAW) facility: Provides more capacity over native FICON channels for programs that process data sets that use striping and compression, such as DB2, VSAM, partitioned dat set extended (PDSE), hierarchical file system (HFS), and z/OS file system (zFS). It does so by reducing channel, director, and control unit processor usage.
- Enhanced problem determination, analysis, and manageability of the storage area network (SAN) by providing registration information to the fabric name server for both FICON and Fibre Channel Protocol (FCP).

> **Note:** The FICON Express8 features support a link data rate of 2, 4, or 8 Gbps. Up to 16 features with up to 64 FICON Express8 channels are supported in two (max) I/O drawers. *FICON Express8 is carry forward only.*

Open Systems Adapter (OSA)

The z13 allows any mix of the supported Open Systems Adapter (OSA) Ethernet features. Up to 48 OSA-Express5S or OSA-Express4S features, with a maximum of 96 ports are supported. OSA-Express5S and OSA-Express4S features are plugged into the PCIe I/O drawer.

> **Note:** The maximum number of combined OSA-Express5S and OSA-Express4S features cannot exceed 48.

OSM and OSX channel-path identifier (CHPID) types

The z13 provides OSA-Express5S, OSA-Express4S CHPID types for Unified Resource Manager (OSM) and for zBX (OSX) connections:

- OSA-Express for Unified Resource Manager (OSM)

  Connectivity to the intranode management network (INMN) is not supported on z13. When the IBM zEnterprise BladeCenter Extension (zBX) Model 002 or Model 003 is upgraded to an IBM z BladeCenter Extension (zBX) Model 004, it becomes an independent node configurable to the ensemble. The zBX model 004 is equipped with two 1U rack-mounted Support elements to manage and control itself and its now independent of the CPC SEs eliminating the need of OSM connectivity.
- OSA-Express for zBX (OSX)

  Connectivity to the intraensemble data network (IEDN) provides a data connection from the z13 to the zBX via OSA-Express5S 10 GbE or OSA-Express4S 10 GbE features.

OSA-Express5S and OSA-Express4S feature highlights

For functions that were previously performed in firmware, the OSA Express5S and OSA-Express4S features now perform those functions in hardware. Additional logic in the IBM application-specific integrated circuit (ASIC) that is included with the feature handles packet construction, inspection, and routing, allowing packets to flow between host memory and the LAN at line speed without firmware intervention. OSA-Express features provide the important benefits for TCP/IP traffic, namely reduced latency and improved throughput for standard and jumbo frames. Performance enhancements are the result of the data router function present in all OSA-Express features.

The z13 supports six types of OSA features:

- OSA-Express5S 10 GbE Long Reach (LR) and Short Reach (SR)
- OSA-Express5S GbE Long Wave (LX) and Short Wave (SX)
- OSA-Express5S Ethernet 1000BASE-T Ethernet
- OSA-Express4S 10 GbE Long Reach and Short Reach (SR): carry forward only
- OSA-Express4S GbE Long Wave and Short Wave (SX): carry forward only
- OSA-Express4S Ethernet 1000BASE-T Ethernet: carry forward only

OSA-Express5S features are a technology refresh of the OSA-Express4S features. For more information about the OSA features, see Chapter 14, "Configure Open System Adaptor (OSA) and RoCE Express" on page 653.

IBM HiperSockets

The IBM HiperSockets™ function is also known as *Internal Queued Direct Input/Output (internal QDIO or iQDIO)*. It is an integrated function of the z13 that provides users with attachments to up to 32 high-speed virtual LANs with minimal system and network processor usage.

HiperSockets can be customized to accommodate varying traffic sizes. Because the HiperSockets function does not use an external network, it can free up system and network resources, eliminating attachment costs while improving availability and performance.

For communications between LPARs in the same z13 server, HiperSockets eliminates the need to use I/O subsystem features and to traverse an external network. Connection to HiperSockets offers significant value in server consolidation by connecting many virtual servers. It can be used instead of certain coupling link configurations in a Parallel Sysplex.

HiperSockets is extended to allow integration with IEDN, which extends the reach of the HiperSockets network outside the CPC to the entire ensemble, and displays it as a single Layer 2 network.

10GbE RoCE Express

The 10GbE RoCE Express feature is an RDMA-capable network interface card. The 10GbE RoCE Express feature is supported on z13, zEC12 and zBC12 and is for use exclusively in the PCIe I/O drawer. Each feature has one PCIe adapter. A maximum of 16 features can be installed.

The 10GbE RoCE Express feature uses a short reach (SR) laser as the optical transceiver, and supports use of a multimode fiber optic cable terminated with an LC Duplex connector. Both point-to-point connection and switched connection with an enterprise-class 10 GbE switch are supported.

Support is provided by z/OS which supports two ports on z13, shared by up to 31 partitions.

Note: zEC12 and zBC12 only supports one port per feature dedicated to one partition.

For more information, see Chapter 14, "Configure Open System Adaptor (OSA) and RoCE Express" on page 653.

1.1.7 Coupling Links support

Support for Parallel Sysplex includes the Coupling Facility Control Code and coupling links. Coupling connectivity in support of Parallel Sysplex environments is provided on the z13 by the following features:

- New PCIe Gen3, Integrated Coupling Adapter (PCIe-O SR) allows 2 ports coupling links connectivity for a distance of up to 150 m (492 feet) at 8GBps each.
- HCA3-O,12x InfiniBand coupling links offer up to 6 GBps of bandwidth between z13, zBC12, z196 and z114 systems, for a distance of up to 150 m (492 feet), with improved service times over previous HCA2-O used in prior z Systems platforms.
- HCA3-O LR, 1x InfiniBand coupling links offer up to 5 Gbps connection bandwidth between z13, zEC12, zBC12, z196 and z114 for a distance of up to 10 km (6.2 miles). The HCA3-O LR (1xIFB) type has twice the number of links per fanout card as compared to type HCA2-O LR (1xIFB) used in the previous z Systems platforms.
- Internal Coupling (ICs) channels operate at memory speed.

All coupling link types can be used to carry Server Time Protocol (STP) messages. The z13 does not support ISC-3 connectivity.

CFCC Level 20

CFCC level 20 is delivered on the z13 with driver level 22. CFCC Level 20 introduces the following enhancements:

- Support for up to 141 ICF processors
  - The maximum number of logical processors in a CF partition remains at 16.
- Large memory support
  - Improves availability and scalability for larger CF cache structures and data sharing performance with larger DB2 Group Buffer Pools (GBP).
  - Removes inhibitors to using large CF structures, enabling use of Large Memory to appropriately scale to larger DB2 Local Buffer Pools (LBP) and Group Buffer Pools (GBP) in data sharing environments.
  - CF structure size remains at a maximum of 1 TB.
- Support for new IBM Integrated Coupling Adapter (ICA)

The z13 systems with CFCC Level 20 require z/OS V1R12 or later, and IBM z/VM® V6R2 or later for guest virtual coupling. z13 (CFCC level 20) can coexist in a sysplex with zEC12 (CFCC levels 19) and z196 (CFCC level 17).

1.1.8 Special purpose features

This section overviews several features that, although installed in the PCIe I/O drawer, provide specialized functions without actually performing I/O operations. That is, no data is moved between the CPC and externally attached devices.

Cryptography

Integrated cryptographic features provide leading cryptographic performance and functionality. Reliability, availability, and serviceability (RAS) support is unmatched in the industry, and the cryptographic solution has received the highest standardized security certification (FIPS 140-2 Level 4[1]). The crypto cards permit you to add or move crypto coprocessors to LPARs without preplanning.

The z13 implements the PKCS #11, one of the industry-accepted standards called public key cryptographic standards (PKCS) provided by RSA Laboratories, the security division of EMC Corporation. It also implements the IBM Common Cryptographic Architecture (CCA) in its cryptographic features.

Configurable Crypto Express5S feature

The Crypto Express5S represents the newest generation of cryptographic feature that is designed to complement the cryptographic capabilities of the CPACF. It is an optional feature of z13 server generation. The Crypto Express5S feature is designed to provide port granularity for increased flexibility with one PCIe adapter per feature. For availability reasons, a minimum of two features are required.

The Crypto Express5S is a state-of-the-art, tamper-sensing, and tamper-responding programmable cryptographic feature that provides a secure cryptographic environment. Each adapter contains a tamper-resistant hardware security module (HSM). The HSM can be configured as a Secure IBM CCA coprocessor, as a Secure IBM Enterprise PKCS #11 (EP11) coprocessor, or as an accelerator:

- Secure IBM CCA coprocessor is for secure key encrypted transactions that use CCA callable services (default).
- Secure IBM Enterprise PKCS #11 (EP11) coprocessor implements an industry standardized set of services that adhere to the PKCS #11 specification v2.20 and more recent amendments.

  This new cryptographic coprocessor mode introduced the PKCS #11 secure key function.
- Accelerator for public key and private key cryptographic operations is used with Secure Sockets Layer/Transport Layer Security (SSL/TLS) acceleration.

Federal Information Processing Standards (FIPS) 140-2 certification is supported only when configured as a CCA or an EP11 coprocessor.

TKE workstation and support for smart card readers

The Trusted Key Entry (TKE) workstation and the TKE 8.0 LIC are optional features on the z13. The TKE workstation offers a security-rich solution for basic local and remote key management. It provides to authorized personnel a method for key identification, exchange, separation, update, backup, and a secure hardware-based key loading for operational and master keys. TKE also provides a secure management of host cryptographic module and host capabilities.

Support for an optional smart card reader attached to the TKE workstation allows the use of smart cards that contain an embedded microprocessor and associated memory for data storage. Access to and the use of confidential data on the smart cards are protected by a user-defined personal identification number (PIN).

> **Note:** When Crypto Express5S is configured as a Secure IBM Enterprise PKCS #11 (EP11) coprocessor, the TKE workstation *is required* to manage the Crypto Express5S feature. If the smart card reader feature is installed in the TKE workstation, the new smart card part 74Y0551 is required for EP11 mode.

For more information about the cryptographic features, see Chapter 10, “Crypto Express5S configuration” on page 541.

[1] Federal Information Processing Standards (FIPS) 140-2 Security Requirements for Cryptographic Modules

Flash Express

The *Flash Express* optional feature is intended to provide performance improvements and better availability for critical business workloads that cannot afford any hits to service levels. Flash Express is easy to configure, requires no special skills, and provides rapid time to value.

Flash Express implements storage-class memory (SCM) in a PCIe card form factor. Each Flash Express card implements an internal NAND Flash solid-state drive (SSD), and has a capacity of 1.4 TB of usable storage. Cards are installed in pairs, which provide mirrored data to ensure a high level of availability and redundancy. A maximum of four pairs of cards (eight features) can be installed on a z13, for a maximum capacity of 5.6 TB of storage. For more information, see Chapter 13, “Flash Express” on page 631.

zEDC Express

zEDC Express, an optional feature that is exclusive to z13, zEC12 and zBC12, provides hardware-based acceleration for data compression and decompression with lower CPU consumption than the previous compression technology on z Systems.

Exploitation of the zEDC Express feature by the z/OS V2R1 zEnterprise Data Compression acceleration capability is designed to deliver an integrated solution to help reduce CPU consumption, optimize performance of compression-related tasks, and enable more efficient use of storage resources, while providing a lower cost of computing and also helping to optimize the cross-platform exchange of data.

One to eight features can be installed on the system. There is one PCIe adapter/compression coprocessor per feature, which implements compression as defined by RFC1951 (DEFLATE). A zEDC Express feature can be shared by up to 15 LPARs.

1.2 Hardware Management Console and Support Element

The Hardware Management Console (HMC) and Support Element (SE) are appliances that together provide platform management for z Systems. When the z13 and zBX Model 004 are members of an ensemble, they are also managed by HMCs and their respective SEs.

In an ensemble, the HMC is used to manage, monitor, and operate one or more zEnterprise CPCs and their associated LPARs, as well as the zBX Model 004s. Also, when the zEnterprise and a zBX Model 004 are members of an ensemble, the HMC has a global (ensemble) management scope, compared to SEs (either on the zBX Model 004 or on the CPCs) which have local (node) management responsibility. When tasks are performed on the HMC, the commands are sent to one or more CPC SEs or zBX Model 004 SEs, which then issue commands to their zEnterprise CPCs and zBXs. To promote high availability, an ensemble configuration requires a pair of HMCs in primary and alternate roles.

For more information, see Chapter 3.2, “Hardware Management Console and Support Element configuration planning” on page 50.

1.3 IBM z BladeCenter Extension (zBX) Model 004

The IBM z BladeCenter Extension (zBX) Model 004 is improving infrastructure reliability by extending the mainframe systems management and service across a set of heterogeneous compute elements in an ensemble.

The zBX Model 004 is available as an optional MES upgrade from a zBX Model 003 or a zBX Model 002, to work with the z13 server and consists of the following components:

- Two internal 1U rack-mounted Support Elements providing zBX monitor and control functions.
- Up to four IBM 42U Enterprise racks.
- Up to eight BladeCenter chassis with up to 14 blades, each with up to two chassis per rack for a maximum of 112[2] blades.
- Intranode management network (INMN) Top of Rack (ToR) switches. On the zBX Model 004, the new local zBX Support Elements directly connect to the INMN within the zBX for management purposes.
- Intraensemble data network (IEDN) ToR switches. The IEDN is used for data paths between the zEnterprise ensemble members and the zBX Model 004, and also for customer data access. The IEDN point-to-point connections use MAC addresses, not IP addresses (Layer 2 connection).
- The 8-Gbps Fibre Channel switch modules for connectivity to a SAN.
- Advanced management modules (AMMs) for monitoring and management functions for all the components in the BladeCenter.
- Power Distribution Units (PDUs) and cooling fans.
- Optional acoustic rear door or optional rear door heat exchanger.

The zBX is configured with redundant hardware infrastructure to provide qualities of service similar to those of z Systems, such as the capability for concurrent upgrades and repairs.

For more information about the zBX Model 4, see Chapter 12, “Preparing IBM z13 for zBX Model 004” on page 595.

Important: The zBX Model 004 uses the blades that were carried forward in an upgrade from a previous model. Customers will be able to install additional entitlements up to the full zBX installed blade capacity if the existing BladeCenter chassis have empty available slots. After the entitlements are acquired from IBM, clients must procure and purchase the additional zBX supported blades to be added, up to the full installed entitlement LIC record, from another source or vendor other than IBM.

[2] The maximum number of blades varies according to the blade type and blade function.

2

Channel subsystem

This chapter presents an overview of the channel subsystem and includes these topics:

- Channel subsystem (CSS)
- Logical partitions (LPARs)
- Subchannel sets
- Channels
- Defining multiple CSSs and the fourth subchannel set
- Channel activation
- CPC drawer: Fanout cards and adapters

2.1 Channel subsystem (CSS)

The channels in the CSS allow the transfer of data between main memory and I/O devices or other platforms under the control of a channel program. Through these channel connections the CSS enables communication from main memory to peripheral devices.

The CSS allows channel I/O operations to continue independently of other operations within the platform. This allows other functions to resume after an I/O operation is initiated.

The following key entities comprise the CSS:

- Channel

  The channel is the communication path from the CSS to the connected control units and I/O devices. The CSS communicates with I/O devices by using channel paths between the CSS and control units.

- Physical channel ID (PCHID)

  A PCHID reflects the physical location of a channel-type interface. A PCHID number is based on the I/O drawer or PCIe I/O drawer location, the channel feature slot number and the port number of the channel feature.

- Channel-path identifier (CHPID)

  A CHPID is a value that is assigned to each channel path of the system that uniquely identifies that path. A total of 256 CHPIDs are supported by one CSS. A CHPID does not directly correspond to a hardware channel port. It is assigned to a PCHID in Hardware Configuration Definition (HCD) or Input/Output Configuration Program (IOCP).

- Subchannel

  A subchannel provides the logical appearance of a device to the program and contains the information that is required for sustaining a single I/O operation. A subchannel is assigned for each device that is defined to the logical partition.

- Control unit (CU)

  A CU provides the logical capabilities necessary to operate and control an I/O device. It adapts the characteristics of each device so that it can respond to the standard form of control that is provided by the CSS. A CU can be housed separately, or it can be physically and logically integrated with the I/O device, the CSS or within the platform itself.

- Input/output (I/O) device

  An I/O device has the characteristics of the peripheral device that it represents. In the simplest case, an I/O device is attached to one control unit and is accessible through one channel path.

Figure 2-1 shows the relationship between the CSS, the channels, the control units, and the I/O devices.

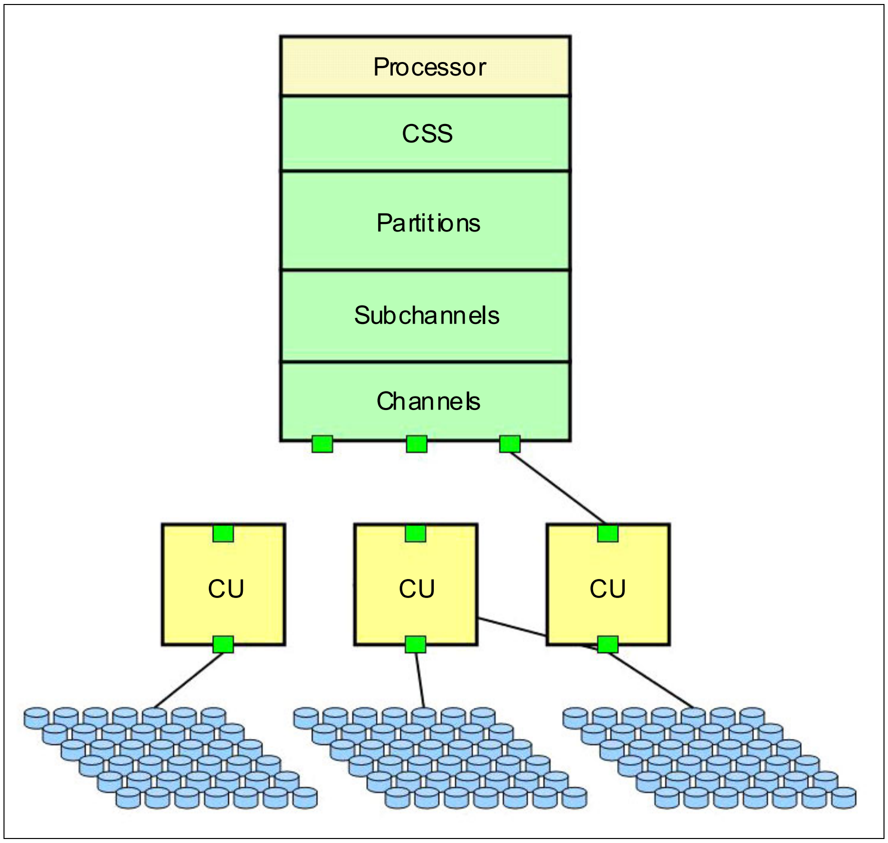

Figure 2-1 Channel subsystem relationship to channels, control units, and I/O devices

2.1.1 Multiple channel subsystems

The multiple CSS concept that is implemented in the z Systems platforms is designed to offer considerable increase in processing power, memory size, and I/O connectivity over previous platforms.

Table 2-1 shows a summary of CSS features available on various z Systems models. A fifth and sixth CSS (CSS4 and CSS5) and a fourth subchannel set (SS3) was introduced in the z13. With SS3, you can define extra 64 K subchannels in each CSS. For more information, see 2.3, “Subchannel sets” on page 20 for details.

Table 2-1 CSS features per z Systems model

Item	**z13**	**zEC12**	**z196**
Number of CSSs	Six per platform	Four per platform	Four per platform
Devices in subchannel set 0	63.75 K per CSS 382.5 K per platform	63.75 K per CSS 255 K per platform	63.75 K per CSS 255 K per platform
Devices in subchannel set 1	64 K: 1 per CSS 384 K: 6 per platform	64 K: 1 per CSS 256 K: 4 per platform	64 K: 1 per CSS 256 K: 4 per platform
Devices in subchannel set 2	64 K: 1 per CSS 384 K: 6 per platform	64 K: 1 per CSS 256 K: 4 per platform	64 K: 1 per CSS 256 K: 4 per platform

Item	z13	zEC12	z196
Devices in subchannel set 3	64 K: 1 per CSS 384K: 6 per platform	N/A	N/A
Partitions	15 per CSS (0 - 4) 10 for CSS5 85 per platform	15 per CSS 60 per platform	15 per CSS 60 per platform
CHPIDs	256 per CSS 1536 per platform	256 per CSS 1024 per platform	256 per CSS 1024 per platform

On a z Systems platform, all channel subsystem images are defined within a single I/O configuration data set (IOCDS). The IOCDS is loaded into the platform's hardware system area (HSA) during a power-on reset (POR).

Figure 2-2 shows a logical view of these relationships. The z13 supports six channel subsystems per platform. CSSs are numbered from 0 to 5. These designations are sometimes referred to as the "CSS Image ID" or "CSSID" (CSSID 0, 1, 2, or 3 for zEC12 and z196) and (CSSID 0, 1, 2, 3, 4 or 5 for z13).

Each channel subsystem can have from 1 to 256 channels and can be configured with 1 to 15 logical partitions (LPARs). See 2.1.2, "Maximum 10 logical partitions in CSS5" on page 18.

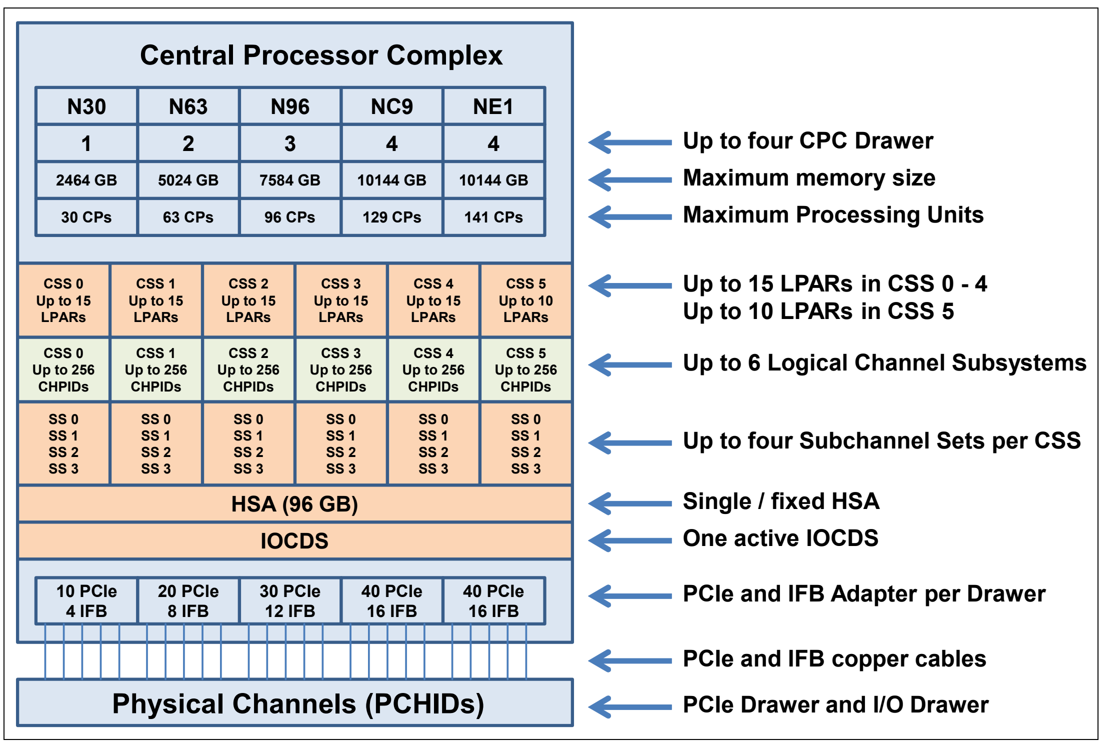

Figure 2-2 Logical view of multiple CSSs in a z13

2.1.2 Maximum 10 logical partitions in CSS5

LPARs B-F in CSS5 will appear in any IOCP deck / IODF as reserved but they are for IBM use only. You receive error messages if you try to add extra logical partition names using those IDs through HCD. Hence the maximum usable LPARs per platform is 85.

2.1.3 Multiple image facility (MIF)

MIF enables resource sharing across LPARs within a single CSS. When a channel is shared across LPARs in multiple CSSs, it is known as *spanning*. See Table 2-2 on page 25 for more information.

2.2 Logical partitions (LPARs)

A LPAR supports the running of an operating system, such as z/OS, and provides central processors (CPs), memory, subchannels, and access to channels.

The following definitions apply to z Systems servers:

- LPAR name

  Defined through HCD or IOCP. It is the name in the RESOURCE statement in the IOCP. The LPAR names must be unique across all CSSs in the server.

- LPAR identifier

  A number in the range of 00 - 7F. It is assigned in the image profile through the Support Element (SE) or the Hardware Management Console (HMC). It is unique across the server and can also be referred to as the User Logical Partition ID (UPID).

  Generally, establish a numbering convention for the LPAR identifiers. Figure 2-3 on page 20 uses the CSS number concatenated to the MIF Image ID which means LPAR ID “3A” is in CSS “3” with MIFID “A”. This convention fits within the allowed range of LPAR IDs and conveys useful information.

- MIF ID

  Defined by using HCD or IOCP in the RESOURCE statement. It is in the range of x'1’ to x'F' and is unique within a CSS but not unique across multiple CSSs. Multiple CSSs can specify the same MIF ID. The MIF ID is also known as Image ID (IID).

Figure 2-3 on page 20 shows the relationship between CSSs, LPARs and their associated MIF IDs.

- TST1, PROD1, and PROD2 are defined to CSS 0.
- TST2, PROD3, and PROD4 are defined to CSS 1.
- TST3 and TST4 are defined to CSS 2.
- PROD5 and PROD6 are defined to CSS 3.
- TST55 and PROD7 are defined to CSS4.
- PROD8 and TST6 are defined to CSS5.

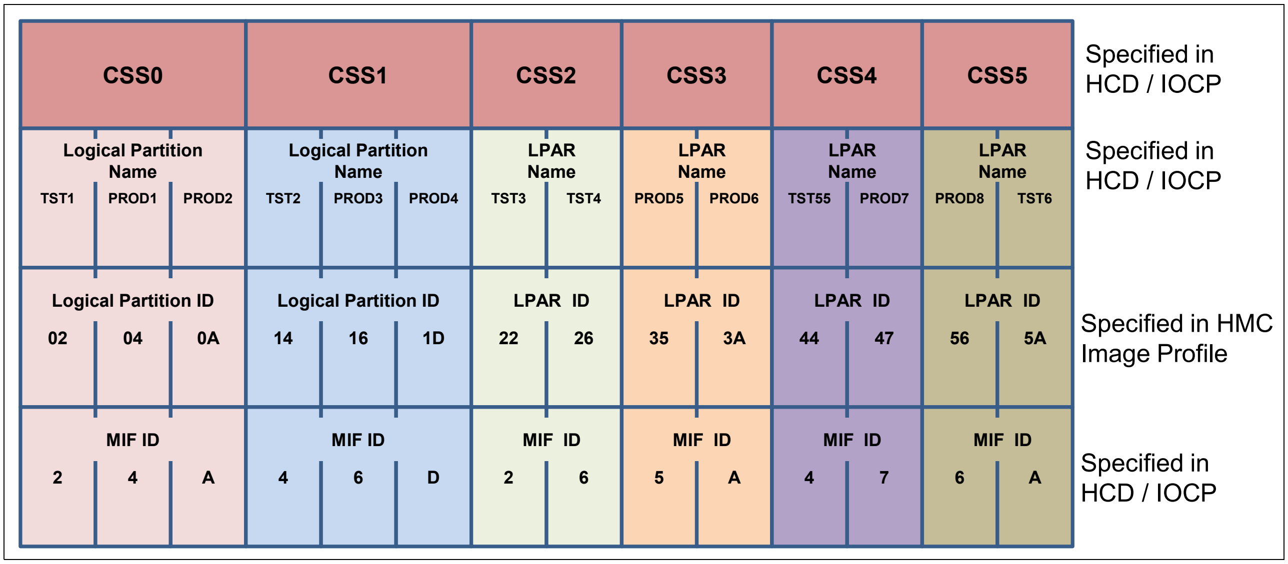

Figure 2-3 CSS and logical partition definition of a z13 server

Notice that PROD1 and TST2 have the same MIFID, although their logical partition IDs differ.

I/O operations for a logical partition are identified as originating from a Channel Image (CI). The Channel Image is defined as CI = CSSID + MIFID.

For IBM FICON I/O operations, the addressing is CI (+ CHPID) + S_ID + D_ID + CUI + UA, where the terms are defined as follows:

S_ID Source ID
D_ID Destination ID
CUI Control Unit ID
UA Unit Address

The FICON control unit logical addressing for channel-to-channel connections is CSSID.MIFID.

2.3 Subchannel sets

In z/Architecture, each I/O device is represented by a separate set of controls for each logical partition that is called a *subchannel*. The subchannel is used by the operating system to pass an I/O request from the System Control Program (SCP) to the channel subsystem. To a program, the subchannel is displayed as a device. In the channel subsystem, the primary control block for an I/O is the *Subchannel*, which is represented by a control block that is called the Unit Control Word (UCW). UCWs are part of the HSA.

The number of devices that can be addressed by a logical partition can be a limitation for some installations. The concept of multiple subchannel sets (MSS) provides relief for this problem.

Usually, a subchannel represents an addressable I/O device. Therefore, a disk control unit with 30 drives uses 30 subchannels for base addresses. An addressable device is associated with a device number and the device number is commonly, but incorrectly, known as the *device address*.

Subchannel numbers are limited to four hexadecimal digits by hardware and software architectures. Four hexadecimal digits provides up to 64-K addresses, which are known as a *set*. IBM reserves 256 subchannels in SS0, leaving 63.75-K subchannels for general use.

For I/O constraint relief, four subchannel sets are available per CSS. This configuration allows access to a greater number of logical volumes. It also allows improved device connectivity for parallel access volumes (PAVs), peer-to-peer remote copy (PPRC) secondaries, and IBM FlashCopy® devices.

The z13 allows you to do an initial program load (IPL) from any subchannel set (SS0-SS3).

The fourth subchannel set was introduced and implemented in each logical channel subsystem by the IBM z13. Subchannel set 0 (SS0) provides a total of 63.75 K subchannels, Subchannel set 1 (SS1), subchannel set 2 (SS2) and subchannel set 3(SS3) provide to the installation the full range of 64 K-1 addresses.

Figure 2-4 shows the CSSs and the associated subchannel sets.

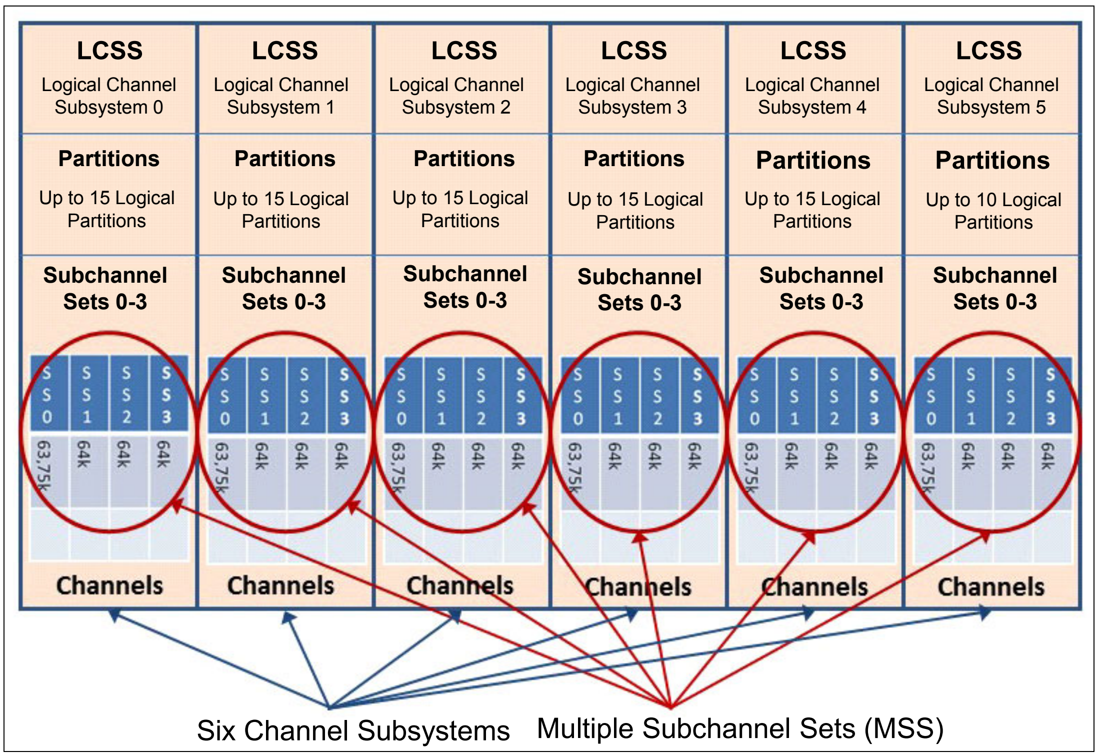

Figure 2-4 Multiple subchannel sets

The current implementation in z/OS allows the use of Subchannel set 0 (SS0) to define any type of device allowed today as base, alias, and secondary devices. The second, third and fourth subchannel sets (SS1-SS3) allow you to define disk alias devices (of both primary and secondary devices), and Metro Mirror secondary devices only.

The appropriate subchannel set number must be included in IOCP definitions or in the HCD definitions that produce the IOCDS. The subchannel set number defaults to zero.

With the availability of HyperPAV, the requirement for PAV devices is greatly reduced. HyperPAV allows an alias address to be used to access any base on the same control unit image per I/O base. It also allows different HyperPAV hosts to use one alias to access different basis, which reduces the number of alias addresses required. HyperPAV is designed to enable applications to achieve equal or better performance than is possible with the original PAV feature alone, while also using the same or fewer z/OS resources. HyperPAV is an optional feature on the IBM System Storage® DS8000® series.

Figure 2-6 on page 26 and Example 2-1 on page 27 illustrate the use and definition of multiple subchannel sets.

2.4 Channels

The CSS communicates with I/O devices through channel paths between the CSS and CUs. A *channel* is the communication path from the CSS to the connected CUs and I/O devices.

2.4.1 Physical Channel ID (PCHID)

A PCHID reflects the physical location of a channel-type interface. A PCHID number is based on the I/O drawer, or PCIe I/O drawer location, the channel feature slot number and the port number of the channel feature. The physical channel, which uniquely identifies a connector jack on a channel feature, is known by its PCHID number. They follow the numbering scheme that is defined for the z13 model. PCHID values are shown in Table 2-7 on page 36 for I/O drawers, and in Table 2-9 on page 39 for PCIe I/O Drawers.

2.4.2 Channel Path ID (CHPID)

A CHPID is a value that is assigned to each channel path of the system that uniquely identifies that path. A total of 256 CHPIDs are supported by one CSS. This provides a maximum of 1536 CHPIDs when six CSSs are defined. Each CHPID within a CSS is associated with a single channel. A CHPID does not directly correspond to a hardware channel port. It is assigned to a PCHID in Hardware Configuration Definition (HCD) or Input/Output Configuration Program (IOCP).

2.4.3 Control unit (CU)

A *control unit* (CU) provides the capabilities necessary to operate and control an I/O device. The CU acts as an interface between the CSS and the I/O device, masking the uniqueness of the I/O device from the CSS. A CU can be housed separately, or it can be physically and logically integrated with the I/O device, the CSS, or within the server itself.

2.4.4 I/O device

An *input/output* (I/O) device is the end point in the “conduit” between a server and a peripheral device. The channel does not communicate directly with I/O devices (it communicates with CUs) and subchannels are displayed as I/O devices to programs.

An I/O device has the characteristics of the peripheral device that it represents. In the simplest case, an I/O device is attached to one control unit and is accessible through one channel path.

2.4.5 Adapter ID (AID)

The AID specifies the adapter identifier that is associated with an InfiniBand Host Channel Adapter (HCA) or an Integrated Coupling Adapter (ICA) SR (PCIe-O SR) used for coupling links. It is determined from the PCHID report when a HCA or ICA is ordered. It reflects the initial physical location of the fanout card. It is bound to the serial number of the card so if the fanout is physically moved the AID moves with it.

For HCA cards, the AID is a two-digit hexadecimal number in the range of x”00” to x”0F”. A maximum of 16 CIB CHPIDs can be defined for the same AID across both ports.

For ICA cards, the AID is a two-digit hexadecimal number in the range of x”10” to x”37”. A maximum of 16 PCIE-O SR adapters are supported per z13. A maximum of 8 CS5 CHPIDs (4 CHPIDs per port) can be defined for the same AID

For more information about InfiniBand coupling links, see *IBM z Systems Connectivity Handbook*, SG24-5444.

2.4.6 Channel path sharing

Two possibilities for channel path sharing are available:

- MIF

  This configuration enables channel sharing among LPARs that run in one channel subsystem.

- Spanning

  This configuration extends the MIF concept of sharing channels across LPARs in a single CSS to sharing channels across LPARs *and* multiple CSSs.

  Spanning is the ability for the channel to be configured to multiple CSSs. When defined that way, the channels can be transparently shared by any or all of the configured LPARs, regardless of the CSS to which they are configured.

MIF-shared channels

IBM Processor Resource/Systems Manager (PR/SM) allows sharing of resources across LPARs. MIF enables channel sharing among LPARs, but sharing is limited to partitions defined to one CSS.

Understanding qualifiers that apply to a LPAR definition is important. The following definitions for IBM z servers are described in 2.2, “Logical partitions (LPARs)” on page 19:

- LPAR name
- LPAR identifier
- MIF ID

Spanned channels

Spanning is the ability of channels to be configured to multiple CSSs, and be transparently shared by any or all of the configured LPARs that are configured in these CSSs. Figure 2-5 shows an example with two CSSs that use spanned channels. There are spanned external channels (with associated PCHIDs) and spanned internal channels (no associated PCHIDs).

Note: Certain channels that are commonly used on earlier servers cannot be spanned. For example, on a z196, IBM ESCON channels can be MIF-shared within one CSS, but cannot be spanned across multiple CSSs.

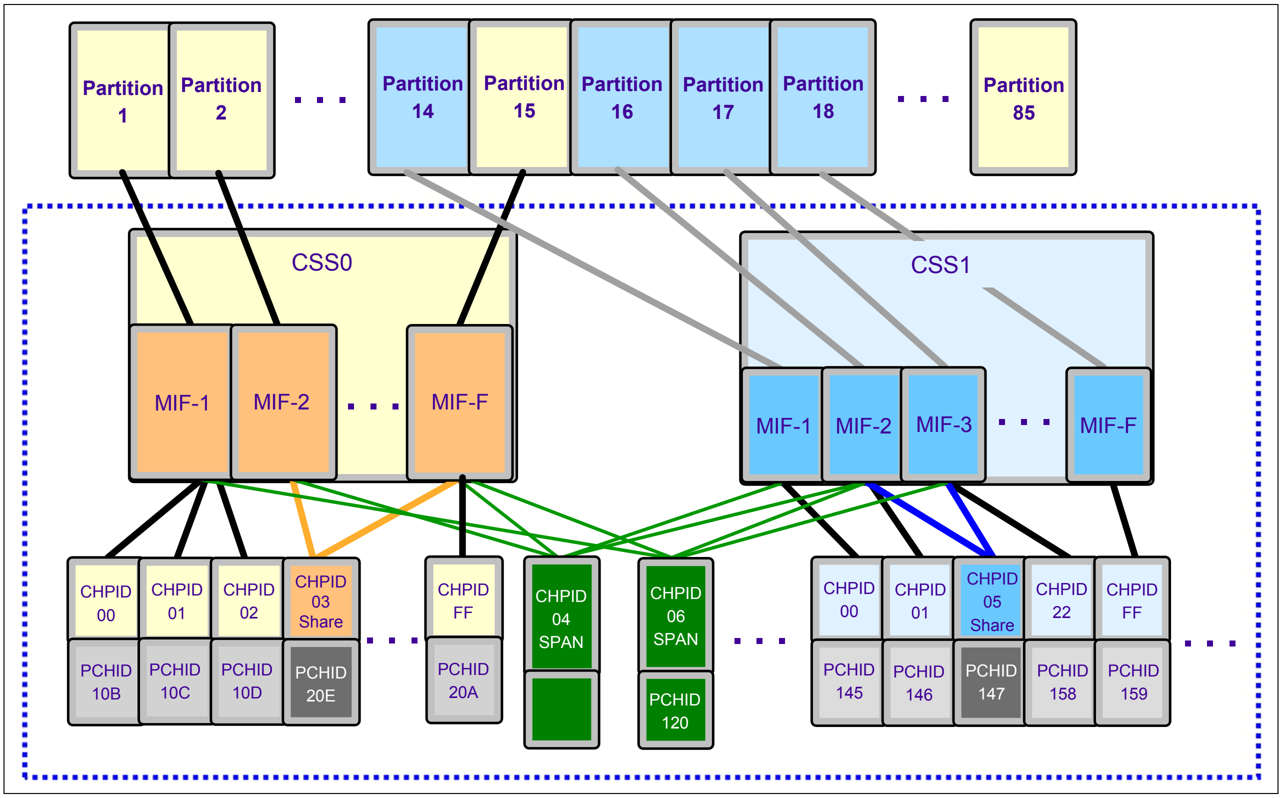

Figure 2-5 MIF-shared and spanned channels

Table 2-2 shows channels that can be shared and spanned on a z13. The channel types are described in more detail in Chapter 3, “Configuration planning” on page 41.

Table 2-2 Spanned and shared channels on a z13 server

Channel type		**CHPID definition**	**MIF-shared channels**	**Spanned channels**
FICON Express16S	External	FC and FCP	Yes	Yes
FICON Express8S[a]	External	FC and FCP	Yes	Yes
FICON Express8[b]	External	FC and FCP	Yes	Yes
OSA-Express5S[a]	External	OSD, OSE, OSC, OSN, OSM, and OSX[c]	Yes	Yes
OSA-Express4S[b]	External	OSD, OSE, OSC, OSN, OSM, and OSX[c]	Yes	Yes
PSIFB	External	CIB	Yes	Yes
ICA	External	CS5	Yes	Yes
IC	Internal	ICP	Yes	Yes
IBM HiperSockets	Internal	IQD	Yes	Yes

a. This feature is available for new-build and carry-forward.
b. This feature is carry-forward only.
c. Depending on feature code.

Figure 2-6 shows a server with six CSSs using a fourth subchannel set.

Note: To fit within the page width, only CSSO, CSS1, CSS5 are shown in this figure.

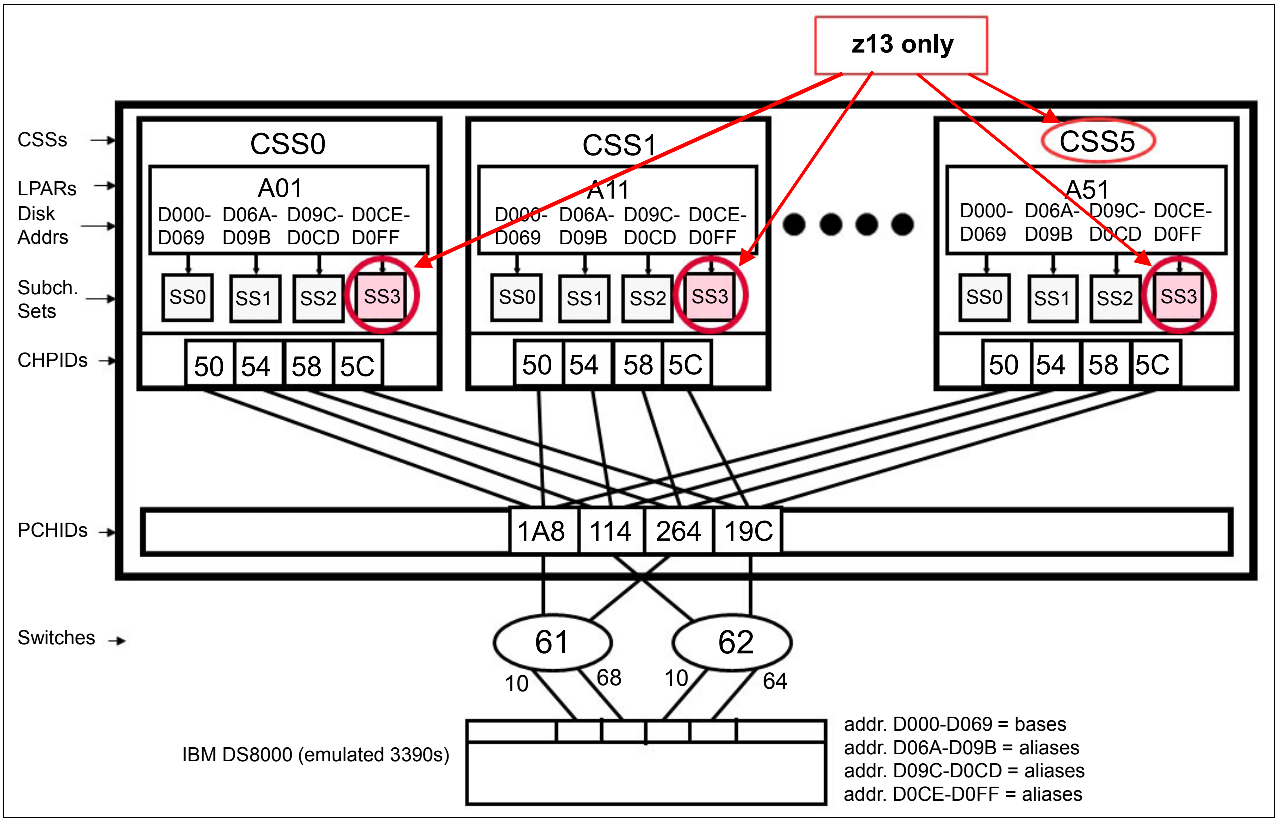

Figure 2-6 Configuration example with CSSs, MSSs, and PCHIDs

- The logical CHPID numbers 50, 54, 58, and 5C are shared channels between the LPARs. These CHPIDs are defined as spanned channel for each physical FICON (TYPE=FC) channels, PCHID numbers 1A8, 114, 264, and 19C.
- The four subchannel sets are used to address the same DASD:
  - SS0 is used to address the base addresses D000-D069,
  - SS1 is used to address the alias addresses D06A-D09B.
  - SS2 is used to address the alias addresses D09C-D0CD.
  - SS3 is used to address the alias addresses D0CE-D0FF.

Example 2-1 on page 27shows an example of IOCDS definitions for six CSSs and four subchannel sets.

The IOCP statements in Example 2-1 on page 27 are not intended to represent a real server, but to illustrate the new elements involved. In the past, dummy LPARs needed to be defined in the IOCP to enable the user to add them dynamically later. Now IOCP automatically reserves all available CSSs and LPARs that were not explicitly defined.

Example 2-1 Multiple CSSs and subchannel sets IOCP definition example

```
ID    MSG1='IODFA3',MSG2='SYS6.IODFA3 - 2014-11-11 11:58',    *
      SYSTEM=(2964,1),                                          *
      TOK=('SCZP501',00800001DA8729641158222601143l5F00000000,*
      00000000,'14-11-11','11:58:22','SYS6','IODFA3')
RESOURCE PARTITION=((CSS(0),(A01,1),(*,2),(*,3),(*,4),(*,5),(**
      ,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F*
      )),(CSS(1),(A11,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*
      *,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(2),*
      (A21,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),*
      (*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(3),(A31,1),(*,*
      2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B)*
      ,(*,C),(*,D),(*,E),(*,F)),(CSS(4),(A41,1),(*,2),(*,3),(**
      ,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D*
      ),(*,E),(*,F)),(CSS(5),(A51,1),(*,2),(*,3),(*,4),(*,5),(*
      *,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,*
      F)))
CHPID PATH=(CSS(0,1,2,3,4,5),50),SHARED,SWITCH=61,PCHID=1A8,  *
      TYPE=FC
CHPID PATH=(CSS(0,1,2,3,4,5),54),SHARED,SWITCH=62,PCHID=114,  *
      TYPE=FC
CHPID PATH=(CSS(0,1,2,3,4,5),58),SHARED,SWITCH=61,PCHID=264,  *
      TYPE=FC
CHPID PATH=(CSS(0,1,2,3,4,5),5C),SHARED,SWITCH=62,PCHID=19C,  *
      TYPE=FC
CNTLUNIT CUNUMBR=D000,                                          *
      PATH=((CSS(0),50,54,58,5C),(CSS(1),50,54,58,5C),(CSS(2),*
      50,54,58,5C),(CSS(3),50,54,58,5C),(CSS(4),50,54,58,5C),(*
      CSS(5),50,54,58,5C)),UNITADD=((00,256)),                  *
      LINK=((CSS(0),10,10,68,64),(CSS(1),10,10,68,64),(CSS(2),*
      10,10,68,64),(CSS(3),10,10,68,64),(CSS(4),10,10,68,64),(*
      CSS(5),10,10,68,64)),CUADD=0,UNIT=2107
IODEVICE ADDRESS=(D000,106),CUNUMBR=(D000),STADET=Y,          *
      PARTITION=((CSS(0),A01),(CSS(1),A11),(CSS(2),A21),(CSS(3*
      ),A31),(CSS(4),A41),(CSS(5),A51)),UNIT=3390B
IODEVICE ADDRESS=(D06A,050),CUNUMBR=(D000),STADET=Y,          *
      PARTITION=((CSS(0),A01),(CSS(1),A11),(CSS(2),A21),(CSS(3*
      ),A31),(CSS(4),A41),(CSS(5),A51)),SCHSET=1,UNIT=3390A
IODEVICE ADDRESS=(D09C,050),CUNUMBR=(D000),STADET=Y,          *
      PARTITION=((CSS(0),A01),(CSS(1),A11),(CSS(2),A21),(CSS(3*
      ),A31),(CSS(4),A41),(CSS(5),A51)),SCHSET=2,UNIT=3390A
IODEVICE ADDRESS=(D0CE,050),CUNUMBR=(D000),STADET=Y,          *
      PARTITION=((CSS(0),A01),(CSS(1),A11),(CSS(2),A21),(CSS(3*
      ),A31),(CSS(4),A41),(CSS(5),A51)),SCHSET=3,UNIT=3390A
```

Note: LPARs B - F in CSS5 show as reserved and available, but see 2.1.2, “Maximum 10 logical partitions in CSS5” on page 18

The SCHSET parameter in the IODEVICE statement defines the subchannel set used for a specific address range. If no subchannel set is defined, it defaults to subchannel set 0. The following devices are defined in the IOCDS:

- Bases D000-D069 CSS0-5 SS0
- Aliases D06A-D09B CSS0-5 SS1
- Aliases D09C-D0CD CSS0-5 SS2
- Aliases D0CE-D0FF CSS0-5 SS3

2.4.7 Channel program

A channel program is a set of channel command words (CCWs). Channel programs are built by the requester of the I/O and then control is passed to the next phase of the I/O which is run by the system assist processor.

2.4.8 System assist processor (SAP)

A system assist processor is a special-purpose processor unit (PU) responsible for handling I/O. System assist processors are sometimes referred to as I/O processors (IOPs).

2.4.9 Hardware system area (HSA)

The HSA is an area of memory in the processor main storage that is used by the hardware. It is established during power-on reset (POR) using the configuration information from the IOCDS. The z13 provides a fixed amount of storage that is reserved for HSA and is fenced off from client purchased memory. The HSA contains the subchannels (UCWs). Table 2-3 shows the HSA size for the z13, zEC12, z196, and z114 servers.

Table 2-3   Fixed HSA sizes

z Systems platforms	Fixed HSA size
IBM z13	96 GB
IBM zEnterprise EC12	32 GB
IBM zEnterprise 196	16 GB
IBM zEnterprise 114	8 GB

2.5 Defining multiple CSSs and the fourth subchannel set

This section explains the new IODF definition for the z13 processor. For more information, see *Hardware Configuration Definition Planning,* GA32-0907 and *Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037.

When the definition of a new processor of type 2964 is added by HCD, the maximum number of logical partitions is automatically generated in input/output definition file (IODF) as shown in Figure 2-7.

```
------------------------------------------------------------------------------
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
****************************** Bottom of data ********************************



+------------------------------------------------------------------+
| Definition of processor SCZP501 has been extended to its maximum |
| configuration.                                                   |
+------------------------------------------------------------------+
```

Figure 2-7 HCD: Adding a z13 processor

Figure 2-8 on page 30 shows an example channel subsystem list for a 2964 (z13). The z13 supports the maximum of six CSSs in a processor and four subchannel sets for each CSS, with the maximum number of devices in each subchannel set as follows:

- The maximum number of devices is 65280 per CSS for Subchannel set 0.
- The maximum number of devices is 65535 per CSS for Subchannel set 1.
- The maximum number of devices is 65535 per CSS for Subchannel set 2.

The maximum number of devices is 65535 per CSS for Subchannel set 3.

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

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

Processor ID . . . : SCZP501      Sphinx

  CSS Devices in SS0    Devices in SS1    Devices in SS2    Devices in SS3
/ ID  Maximum + Actual  Maximum + Actual  Maximum + Actual  Maximum + Actual
_ 0   65280     106     65535     50      65535     50      65535     50
_ 1   65280     106     65535     50      65535     50      65535     50
_ 2   65280     106     65535     50      65535     50      65535     50
_ 3   65280     106     65535     50      65535     50      65535     50
_ 4   65280     106     65535     50      65535     50      65535     50
_ 5   65280     106     65535     50      65535     50      65535     50
```

Figure 2-8 Example channel subsystem list for z13 processor

Figure 2-9 shows a sample partition list in a channel subsystem. Initially, all 15 LPARs supported by a CSS are assigned as “reserved.” This means that all numbers (between x"1" and x"F") of the partition in each channel subsystem are already defined. This setting is indicated by an asterisk (*) in the partition name field. The asterisk ensures that the CSS treats this logical partition as a reserved logical partition. This configuration allows for the dynamic activation of an LPAR.

```
 Row 1 of 15
Command ===> ________________________________________ Scroll ===> CSR

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

Processor ID  . . . . : SCZP501   Sphinx
Configuration mode  . : LPAR
Channel Subsystem ID  : 0

/ Partition Name   Number Usage + Description
_ A01              1      CF/OS   A01
_ *                2      CF/OS   ________________________________
_ *                3      CF/OS   ________________________________
_ *                4      CF/OS   ________________________________
_ *                5      CF/OS   ________________________________
_ *                6      CF/OS   ________________________________
_ *                7      CF/OS   ________________________________
_ *                8      CF/OS   ________________________________
_ *                9      CF/OS   ________________________________
_ *                A      CF/OS   ________________________________
_ *                B      CF/OS   ________________________________
_ *                C      CF/OS   ________________________________
_ *                D      CF/OS   ________________________________
_ *                E      CF/OS   ________________________________
_ *                F      CF/OS   ________________________________
```

Figure 2-9 Example partition list for a z13 processor

Note: Figure 2-9 shows all 15 LPARs on one screen for ease of reference in this book; the actual panel displays the list over two screens

The partition usage is also defined.

CF/OS means that these reserved logical partitions can be defined as operating system, coupling facility, or Linux logical partition.

The reserved logical partitions can be seen only as definitions in HCD or in the IOCP deck that is produced by HCD. The Resource Partition statement in IOCP shows the reserved partitions with asterisk (‘*’) in the name field. Space is allocated in the HSA. The reserved logical partitions will not be visible in the list of LPARs for the CPC in the Hardware Management Console (HMC).

In the IOCP deck that is produced by HCD, the RESOURCE PARTITION statement shows the reserved partitions with ‘*’ in the name field as shown in Example 2-2.

Example 2-2 Sample from IOCP deck showing just RESOURCE PARTITION statement

```
ID    MSG1='IODFA3',MSG2='SYS6.IODFA3 - 2014-11-11 11:58',    *
      SYSTEM=(2964,1),                                        *
      TOK=('SCZP501',00800001DA8729641158222601143l5F00000000,*
      00000000,'14-11-11','11:58:22','SYS6','IODFA3')
RESOURCE PARTITION=((CSS(0),(A01,1),(*,2),(*,3),(*,4),(*,5),(**
      ,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F*
      )),(CSS(1),(A11,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*
      *,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(2),*
      (A21,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),*
      (*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(3),(A31,1),(*,*
      2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B)*
      ,(*,C),(*,D),(*,E),(*,F)),(CSS(4),(A41,1),(*,2),(*,3),(**
      ,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D*
      ),(*,E),(*,F)),(CSS(5),(A51,1),(*,2),(*,3),(*,4),(*,5),(*
      *,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,*
      F)))
```

For more information about how to manage reserved logical partitions (including dynamic addition and deletion of LPARs) see Chapter 11, “Logical partitions” on page 573.

2.6 Channel activation

No single component or single part of the CPC completely defines the channel subsystem on the z13. It comprises many CPC resources (both hardware and IBM Licensed Internal Code (LIC)). These resources work together to support I/O operations across the CPC. These operations include I/O queuing, de-queuing, priority management, and identification of all I/O operations that are performed by LPARs.

2.6.1 Activation processes

During initialization of the server the CSS rules are enforced. Initialization of the server includes definition and activation of the LPARs. The processes that occur during activation are as follows:

1. Perform power-on-reset (POR) to activate the z13 CPC (Reset Profile). The POR will perform the following steps:
   a. Build the CSS HSA contents based on the I/O configuration definition in the selected IOCDS.
   b. Initialize all defined channel types.
   c. Initialize FICON links.
2. Activate the required LPARs (image profile). LPAR activation will perform the following steps:
   a. Initialize the LPAR as per the LPAR image profile.
   b. Assign storage to the LPAR. The storage is never shared with other LPARs
   c. For FICON channels, establish logical paths.
3. Perform the initial program load (IPL; load profile or manual load). An IPL involves the following steps:
   a. Perform an I/O system reset for the logical partition for all defined channel paths.
   b. The operating system starts the required I/O operations.

2.7 CPC drawer: Fanout cards and adapters

A fanout card provides connectivity from the Central Processor Complex (CPC) to the I/O drawers or between CPCs (coupling links).

The z13 server can have one to four CPC drawers installed. The PCIe Gen 3 and InfiniBand (IFB) fanouts are in the front of each drawer. Each drawer has 14 fanout slots (10 PCIe Gen3 and 4 IFB).

Figure 2-10 shows the layout of a z13 CPC drawer.

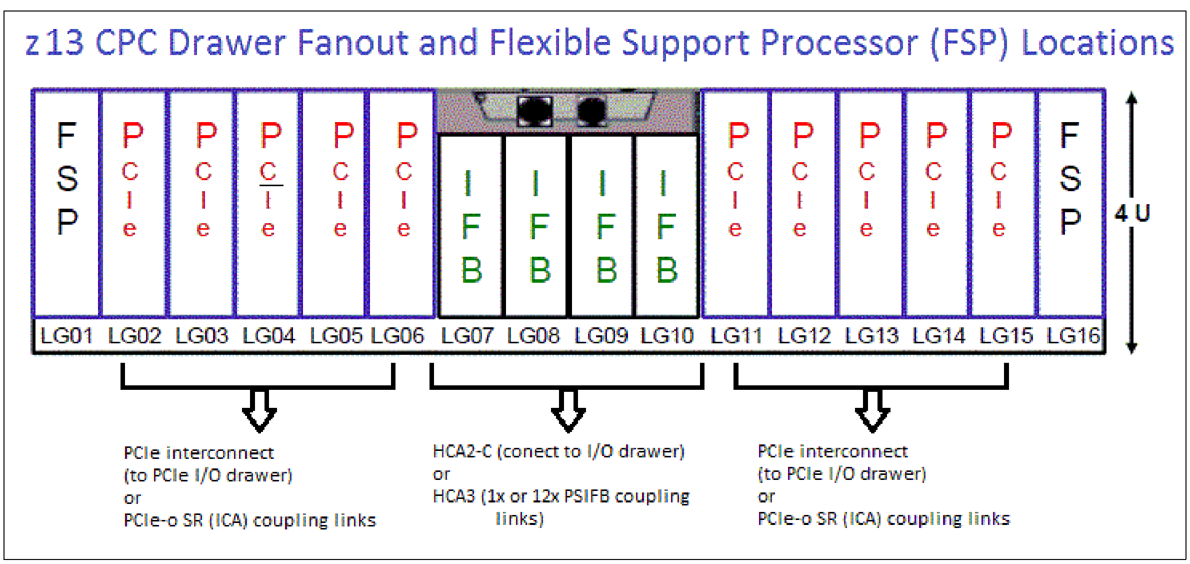

Figure 2-10 z13 CPC drawer fanout and Flexible Support Processor (FSP) locations

The fanout cards in the following locations are supported in the CPC drawer by the z13:

- LG02-LG06, LG11-LG15
  - PCIe Gen3 Interconnect

    This copper fanout provides connectivity to the *PCIe I/O drawer*.
  - Integrated Coupling Adapter (ICA) PCIe-O Gen3 SR

    This optical fanout provides short range coupling link connectivity up to 150 meters distance to another ICA PCIe-O Gen3 SR.
- LG07-LG10
  - Host Channel Adapter (HCA2-C)

    This copper fanout provides connectivity to the *I/O drawer*.

    Note: HCA2-C can only be carried forward on an upgrade to a z13.
  - Host Channel Adapter (HCA3-O (12xIFB))

    This optical fanout provides 12x InfiniBand coupling link connectivity up to 150 meters distance to an HCA3-O on a z13, zEC12, z196 or z114.
  - Host Channel Adapter (HCA3-O LR (1xIFB))

    This optical long range fanout provides 1x InfiniBand coupling link connectivity up to 10 km unrepeated distance to an HCA3-O LR on a z13, zEC12, z196 or z114.

Fanouts that are used for coupling links have an assigned AID. This AID must be used for definitions in IOCDS to have a relation between the physical adapter location and the CHPID number. See 2.4.5, “Adapter ID (AID)” on page 23 and for AID number assignment, see “Adapter ID (AID) and port assignment” on page 111.

2.7.1 I/O drawer and PCIe I/O drawer

I/O drawer, PCIe I/O drawer and their combinations are described in this section.

I/O drawer

I/O drawers can only be carried forward as part of an upgrade to the z13.

I/O drawer carried forward notes:

- The only features in an I/O drawer are FICON Express8.
- Any I/O drawer limits maximum memory or storage in any LPAR in the system to 1TB
- The total fanout slots for PSIFB will be reduced by 2 (independent of whether one or two I/O drawers are carried forward - there are two HCA2-C fanout cards required in both cases).
- The maximum number of PCIe I/O features will be reduced by 32.
- The number of maximum FICON channel will be reduced by 32.

A maximum of two I/O drawers for carry forward of FICON Express8 cards is supported. This represents a maximum of 16 FICON Express8 cards (a total of 64 FICON Express8 ports).

PCIe I/O drawer

For a *new build* of a z13, a maximum of five PCIe I/O drawers can be installed supporting up to 160 PCIe I/O features.

PCIe I/O drawers support the following features:

- FICON Express16S
- FICON Express8S
- OSA Express5S
- OSA Express4S (Carry forward only)
- 10GbE RoCE Express
- Crypto Express5S
- Flash Express
- zEDC Express

I/O drawer and PCIe I/O drawer combinations

A combination of *I/O drawers* and *PCIe I/O drawers* is only available on upgrades (carry forward). The z13 supports the following I/O drawer and PCIe I/O drawers combinations:

- No I/O drawers and a maximum of five PCIe I/O drawers
- One I/O drawer and a maximum of four PCIe I/O drawers
- Two I/O drawer and a maximum of three PCIe I/O drawers

Table 2-4 gives an overview of the number of I/O drawers and PCIe I/O drawers that can be present in a z13.

Table 2-4 I/O drawer and PCIE I/O drawer summary

Description	New build	Carry forward	MES add
I/O Drawer	0	0-2	0[a]
PCIe I/O Drawer	0-5	0-5	0-5

a. Empty slots in a carried forward I/O drawer CANNOT be filled by a MES

The number of I/O drawers, or PCIe I/O drawers can be driven by the total number of cards in a configuration, or by the drawer power budget.

A maximum of 16 I/O features (FICON Express8) can be carried forward. Table 2-5 lists the number and mix of features in a combination of I/O drawers and PCIe I/O drawers.

Table 2-5 Number and mix of features in I/O and PCIe I/O drawers

Carried forward (FICON Express8 only) features	Number of 8 slot I/O drawers required	Maximum remaining PCIe I/O drawers / slots available for other “I/O” features
0	0	5 / 160
1-8	1	4 / 128
9-16	2	3 / 96

Notes: On *new builds* of the z13 only PCIe I/O drawers are supported.

2.7.2 I/O drawer domains, slots and PCHIDs

Each I/O drawer supports two I/O domains and a total of eight I/O cards (four front, four rear) with 4 ports per card. Each I/O domain supports four I/O cards.

Figure 2-11 shows the physical layout of an I/O drawer.

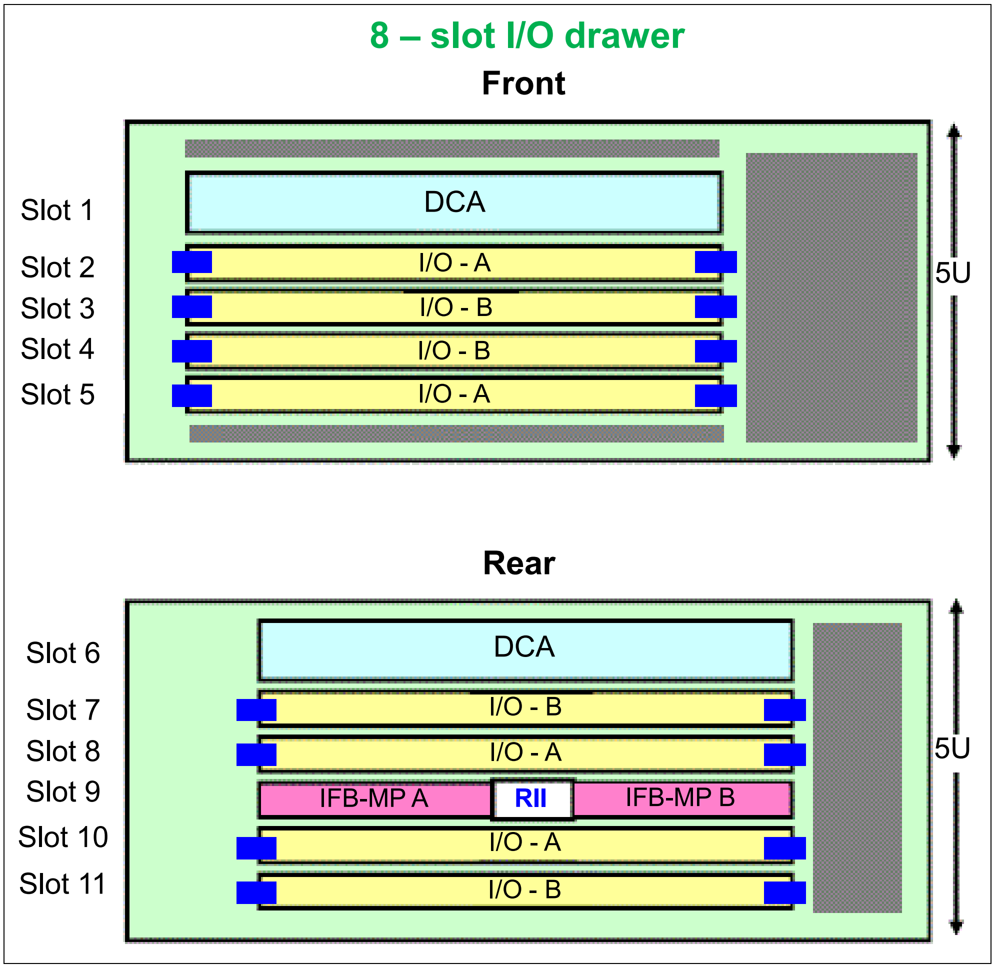

Figure 2-11 I/O drawer physical layout

I/O slot 9 contains the IFB-MP card that is used to connect to the CPC drawer (via the HCA2-C fanout). Two connections are always provided to allow redundant I/O interconnect. Figure 2-12 on page 36 shows the I/O drawer domains and the connectivity to them.

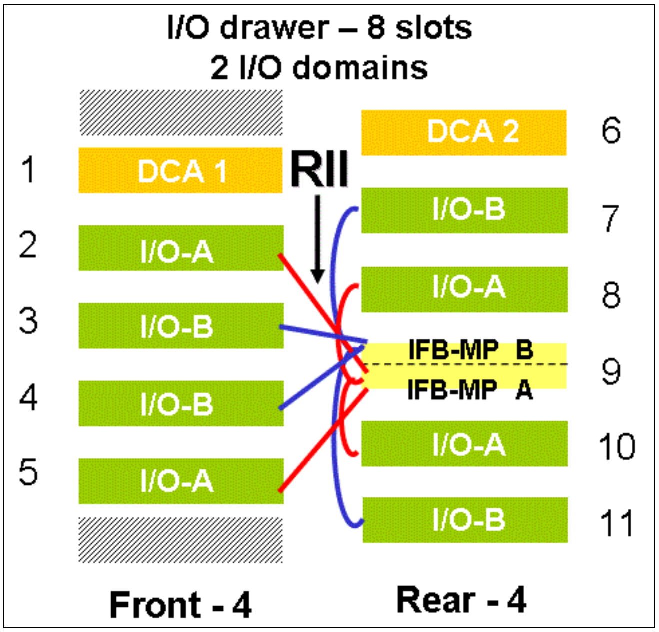

Figure 2-12 I/O drawer domains

The assignment of slots to domains in each I/O drawer is identified in Table 2-6.

Table 2-6 I/O domain assignment in the I/O drawer

I/O domains	I/O slots in domain
A	2, 5, 8, and 10
B	3, 4, 7, and 11

The PCHID number range for each I/O card is determined by the physical location of the I/O card in an I/O drawer. Table 2-7 lists the PCHID number range for each I/O slot in each I/O drawer.

Table 2-7 PCHID assignments for I/O drawers

I/O drawer slot	PCHID range	
	Drawer 1 (Z22B)	Drawer 2 (Z15B)
2	100 - 103[a]	180 - 183[a]
3	110 - 113[a]	190 - 193[a]
4	120 - 123[a]	1A0 - 1A3[a]
5	130 - 133[a]	1B0 - 1B3[a]
7	140 - 143[a]	1C0 - 1C3[a]
8	150 - 153[a]	1D0 - 1D3[a]
10	160 - 163[a]	1E0 - 1E3[a]
11	170 - 173[a]	1F0 - 1F3[a]

a. I/O drawers have only 4 ports per slot and only 8 slots total so the PCHID address ranges used for I/O drawer slots are the first 4 addresses of the ranges that would have been used if a PCIe I/O drawer had been in place

2.7.3 PCIe drawer domains, slots, and PCHIDs

Each PCIe I/O drawer supports up to 4 I/O domains (0,1,2,3) and a total of 32 I/O cards. Each I/O domain supports eight I/O cards.

Figure 2-13 shows the physical layout of a PCIe I/O drawer.

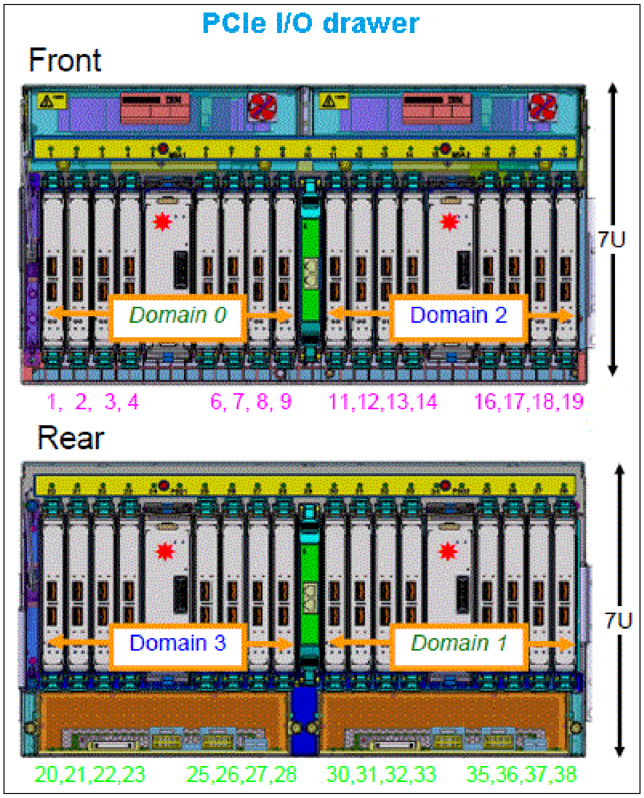

Figure 2-13 PCIe I/O drawer physical layout

The PCIe I/O Drawer supports up to 32 I/O cards. They are organized in four hardware domains per drawer. Each I/O domain supports up to eight cards which can be FICON, OSA, RoCE, Crypto, Flash and zEDC Express feature. Each domain is driven through a PCIe interconnect card. These PCIe interconnect cards occupy slots 5, 15, 24, and 34. Two PCIe interconnect cards always provide a backup path for each other through the passive connection in the PCIe I/O Drawer backplane. During a PCIe fanout card or cable failure, all 16 I/O cards in the two domains can be driven through a single PCIe interconnect card.

The PCIe I/O drawer domains and the connectivity to them is shown in Figure 2-14.

Note: Figure 2-14 is a “vertical” representation of the layout of a PCIe I/O drawer for ease of reading.

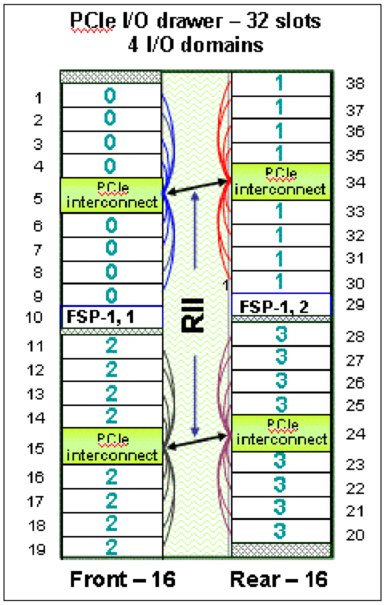

Figure 2-14 PCIe I/O domains

Table 2-8 lists the I/O domains and I/O slots per domain.

Table 2-8 I/O domains of PCIe I/O drawer

Domain	I/O slot in domain
0	01, 02, 03, 04, 06, 07, 08, 09
1	30, 31, 32, 33, 35, 36, 37, 38
2	11, 12, 13, 14, 16, 17, 18, 19
3	20, 21, 22, 23, 25, 26, 27, 28

Table 2-9 on page 39 lists the PCHID assignments for slots in the PCIe I/O drawers. Only the active ports on an installed card are assigned a PCHID. The remainder are unused.

Table 2-9 PCHID assignments for PCIe I/O drawers

PCIe I/O drawer slot	PCHID range				
	Drawer 1 (Z22B)	**Drawer 2 (Z15B)**	**Drawer 3 (Z08B)**	**Drawer 4 (Z01B)**	**Drawer 5 (A32B)**
1	100 - 103	180 - 183	200 - 203	280 - 283	300 - 303
2	104 - 107	184 - 187	204 - 207	284 - 287	304 - 307
3	108 - 10B	188 - 18B	208 - 20B	288 - 28B	308 - 30B
4	10C - 10F	18C - 18F	20C - 20F	28C - 28F	30C - 30F
6	110 - 113	190 - 193	210 - 213	290 - 293	310 - 313
7	114 - 117	194 - 197	214 - 217	294 - 297	314 - 317
8	118 - 11B	198 - 19B	218 - 21B	298 - 29B	318 - 31B
9	11C - 11F	19C - 19F	21C - 21F	29C - 29F	31C - 31F
11	120 - 123	1A0 - 1A3	220 - 223	2A0 - 2A3	320 - 323
12	124 - 127	1A4 - 1A7	224 - 227	2A4 - 2A7	324 - 327
13	128 - 12B	1A8 - 1AB	228 - 22B	2A8 - 2AB	328 - 32B
14	12C - 12F	1AC - 1AF	22C - 22F	2AC - 2AF	32C - 32F
16	130 - 133	1B0 - 1B3	230 - 233	2B0 - 2B3	330 - 333
17	134 - 137	1B4 - 1B7	234 - 237	2B4 - 2B7	334 - 337
18	138 - 13B	1B8 - 1BB	238 - 23B	2B8 - 2BB	338 - 33B
19	13C - 13F	1BC - 1BF	23C - 23F	2BC - 2BF	33C - 33F
20	140 - 143	1C0 - 1C3	240 - 243	2C0 - 2C3	340 - 343
21	144 - 147	1C4 - 1C7	244 - 247	2C4 - 2C7	344 - 347
22	148 - 14B	1C8 - 1CB	248 - 24B	2C8 - 2CB	348 - 34B
23	14C - 14F	1CC - 1CF	24C - 24F	2CC - 2CF	34C - 34F
25	150 - 153	1D0 - 1D3	250 - 253	2D0 - 2D3	350 - 353
26	154 - 157	1D4 - 1D7	254 - 257	2D4 - 2D7	354 - 357
27	158 - 15B	1D8 - 1DB	258 - 25B	2D8 - 2DB	358 - 35B
28	15C - 15F	1DC - 1DF	25C - 25F	2DC - 2DF	35C - 35F
30	160 - 163	1E0 - 1E3	260 - 263	2E0 - 2E3	360 - 363
31	164 - 167	1E4 - 1E7	264 - 267	2E4 - 2E7	364 - 367
32	168 - 16B	1E8 - 1EB	268 - 26B	2E8 - 2EB	368 - 36B
33	16C - 16F	1EC - 1EF	26C - 26F	2EC - 2EF	36C - 36F
35	170 - 173	1F0 - 1F3	270 - 273	2F0 - 2F3	370 - 373
36	174 - 177	1F4 - 1F7	274 - 277	2F4 - 2F7	374 - 377
37	178 - 17B	1F8 - 1FB	278 - 27B	2F8 - 2FB	378 - 37B
38	17C - 17F	1FC - 1FF	27C - 27F	2FC - 2FF	37C - 37F

3

Configuration planning

This chapter addresses the planning steps to take when you install an IBM z13. These steps apply whether you are upgrading from an existing server or installing a new z13.

This chapter includes the following topics:

- Tools
- Hardware Management Console and Support Element configuration planning
- Channel considerations

3.1 Tools

This section summarizes the various tools available for the IBM z Systems platforms. Table 3-1 lists the machine types for the current z Systems platforms.

Table 3-1 Machine types

Server	Abbreviated name	Machine type
IBM z13	z13	2964
IBM zEnterprise EC12	zEC12	2827
IBM zEnterprise 196	z196	2817
IBM zEnterprise 114	z114	2818
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 Hardware Configuration Definition (HCD) and Channel Path Identifier (CHPID) Mapping Tool (CMT), that refer to the machine type as opposed to server names.

3.1.1 Resource Link

The first step in planning for the installation of the z13 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 this registration process.

After you have an IBM ID, you can customize your profile to accommodate the servers that you are responsible for.

After you access Resource Link, you have access to various resources and tools that are needed for the installation process, and information about the z13. A number of tools are available to simplify the installation process of a z13. Even if you have worked with most of them before, be sure to check for the latest versions that are relevant to z13. The tools include the CHPID Mapping Tool (CMT) and the Coupling Facility Structure Sizer Tool (CFSizer).

The Education and Library tab displays information about the IBM z Systems family and some online tutorials. Under the Tools tab, you can download the latest frequently used tools and get system information and configuration.

See the IBM Resource Link web page:

http://www.ibm.com/servers/resourcelink

3.1.2 Hardware Configuration Definition (HCD)

The HCD is a tool 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, as opposed to writing your own IOCP statements. The validation checking that HCD runs as you enter data helps minimize the risk of errors.

HCD support for the z13 is available on all supported z/OS levels with 2964DEVICE PSP and z/VM beginning with Version 6.2. HCD provides the capability to make both dynamic hardware and software I/O configuration changes. For z/OS, HCD is required for Dynamic I/O reconfiguration.

HCD provides enhancements to the I/O Autoconfiguration function that was introduced in z/OS V2R1. In addition to switched FICON connected controllers, I/O Autoconfiguration can now discover FICON direct-attached controllers and devices and proposes point-to-point connection paths if available.

I/O Autoconfiguration supports the inclusion or exclusion of specific switches or CHPIDs into the discovery and proposal process that users can explicitly specify with the invocation of an I/O Autoconfiguration request. For this purpose, HCD introduces four autoconfigurations. The new policy keywords are as follows:

- AUTO_CHPID_INCLUDE
- AUTO_CHPID_EXCLUDE
- AUTO_SWAD_INCLUDE
- AUTO_SWAD_EXCLUDE

The Autoconfiguration policy keyword AUTO_SS_DEVNUM_SCHEME accepts a new value of NONE. This value bypasses control unit and device number proposals by HCM, and lets the user manually apply the numbers for detected objects.

For more information, see Chapter 4, “I/O Autoconfiguration function (zDAC)” on page 117.

New hardware might require the application of program temporary fixes (PTFs) to enable support in HCD.

Get the most updated information about HCD at this web page:

http://www.ibm.com/systems/z/os/zos/features/hcm

3.1.3 Hardware Configuration Manager (HCM)

HCM provides a graphical user interface for HCD and the associated IODF. It can also define and store more information about the physical hardware to which the IODF is defined.

If your installation requires this feature, or you prefer to build and maintain the IODFs by using this interface, you might find that HCM is a useful tool.

HCM does not replace HCD. Instead, 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.

Get the most updated information about HCM at this web page:

http://www.ibm.com/systems/z/os/zos/features/hcm

3.1.4 Input/output configuration program (IOCP)

ICP IOCP V5 is required for a z13 server. You can define the z13 configuration by using only IOCP. However, HCD is suggested because of its error detection and management features. By using ICP IOCP, writing an IOCDS in preparation for a CPC upgrade is possible.

You can write an IOCDS to a zEC12 or z196, regardless of whether the operating system supports z13, by using these methods:

- In z/OS or IBM z/VSE®, specify the CHECKCPC=NO keyword in the PARM field of the EXEC statement.
- From the z/VM Conversation Monitor Program (CMS), use the NOCHKCPC option.
- When you use HCD, tell it to build an IOCDS in anticipation of an upgrade.
- To support the Open System Adapter (OSA) for zBX management network channel paths (TYPE=OSM) and the OSA for zBX data network channel paths (TYPE=OSX), use ICP IOCP Version 3 Release 1.0 or later.

The IOCDS cannot be used to power on reset the CPC until it is upgraded to a z13. Write an IOCDS when you are preparing to upgrade to a z13 from a zEC12 or a z196 to make the IOCDS ready for your new z13. IOCP can write an IOCDS in preparation for an upgrade to a CPC for which IBM does not offer an upgrade path. The IOCDS will be unusable on the earlier CPC.

Stand-alone IOCP

Stand-alone IOCP programs are run in a logical partition (LPAR). When moving from a zEC12 or z196 to a z13 server and running an earlier operating system version that is not supported on a z13 server, create a stand-alone IOCP for the new environment.

One approach is to upgrade to a level of z/OS or z/VM that supports z13 before the upgrade then apply the supporting HCD and IOCP PTFs. After that is complete, follow the guidance in Chapter 5, “Upgrading from an IBM zEC12 to IBM z13” on page 159 about writing an IOCDS in anticipation of an upgrade.

If you are a z/VM user who does not use HCD or you are a z/VSE user, apply the IOCP PTFs that support z13. Use IOCP to verify your z13 IOCP input statements and then write an IOCDS with IOCP by using the write regardless option (CHECKCPC=NO parm).

If these options are not available, the user must have access to an operating system that can have the supporting IOCP PTFs applied. Generate the user’s IOCP statements, verify they are correct by using IOCP (WRTCDS=NO parm), and then transfer the IOCP input statements to a USB flash drive. Provide the flash drive to the installation team so they can use it as an input for running a stand-alone IOCP during the installation.

For more information about the changes and requirements for ICP IOCP, see the *IOCP User’s Guide,* SB10-7037**.**

3.1.5 CHPID Mapping Tool (CMT)

The CHPID Mapping Tool (CMT) provides a mechanism to map physical channel IDs (PCHIDs) to CHPIDs as required on a z13. The CHPID Mapping Tool is optional but is preferred rather than manually mapping the PCHIDs into CHPIDs. Using the CMT provides the best availability recommendations for a particular configuration.

HCD and the CMT

The HCD process flow for a new z13 installation is shown in Figure 3-1 on page 45.

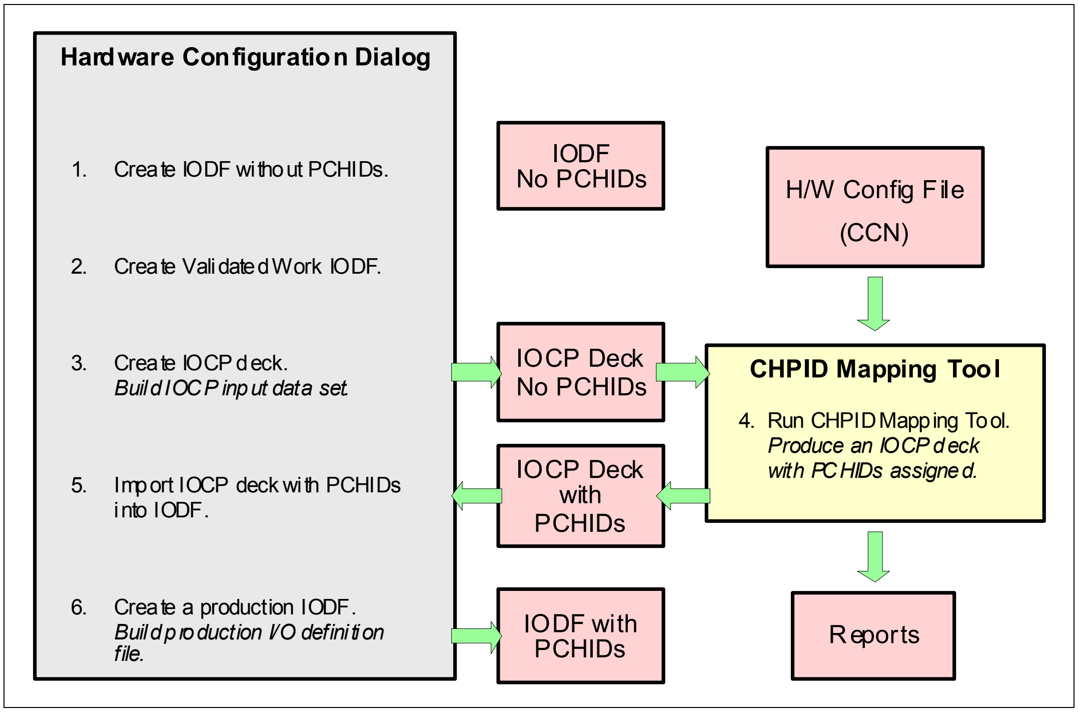

Figure 3-1 Tools: z13 I/O configuration definition flow for a new installation

To download the CMT, log in to the Resource Link site with a registered Resource Link ID:

https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/chpidmain?Opendocument

After the web page opens, click **CHPID Mapping Tool** under the Servers column on the Tools page and complete the following steps:

1. Under Downloads, click **CHPID Mapping Tool**. You can either download the complete CMT program or a file that upgrades an existing CMT program that you might already have installed on your workstation. CHIPID Mapping Tool version 6.17 is the version that is required to support z13 servers.
2. Click the appropriate link and install or upgrade the CHPID Mapping Tool program.

For more information, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

CHPID Mapping Tool support in HCM (without HCD)

HCM for z/OS V2.1 interfaces with the CHPID Mapping Tool. It helps export and import an IOCP and starts the CHPID Mapping Tool. It also creates an IOCP statements file and transfers the file to your workstation and from your workstation back to the host.

You must be logged on to HCM and using the "work" IODF that contains the server that requires processing by the CMT. You also must have the CHPID Mapping Tool installed on the same workstation.

To use the tool, click **Utilities** → **CHPID Mapping Tool Support**. Figure 3-2 on page 46 shows the options that you can select:

- Export IOCP File for the CHPID Mapping Tool (CMT)
- Launch the CHPID Mapping Tool (CMT)
- Import IOCP File from the CHPID Mapping Tool (CMT)

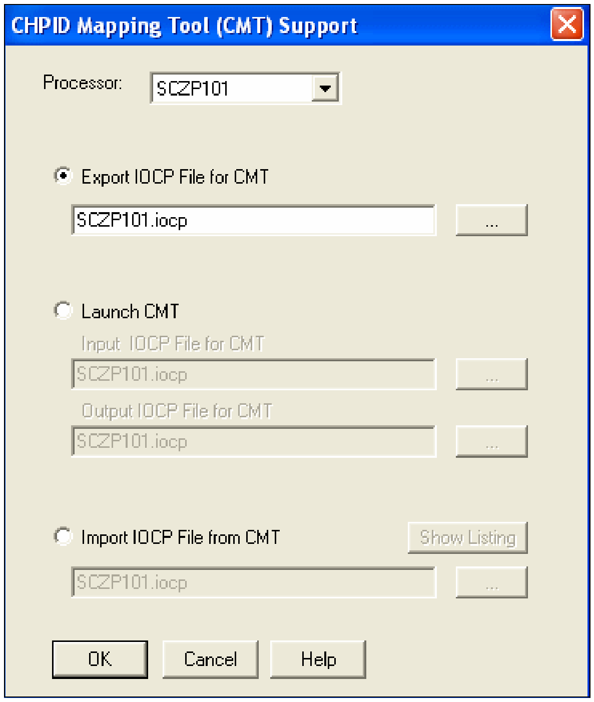

Figure 3-2 CHPID Mapping Tool Support

> **Remember:** Always check the IBM Resource Link website to verify that you have the latest version of the CHPID Mapping Tool installed:
>
> https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/chpidmain?Opendocument

I/O configuration data (Customer Control Number, CCN)

The Customer Control Number (CCN) is generated by your IBM Client Representative when building your configuration order. This number is entered into Resource Link to download a CFReport file which is used as input into the CHPID Mapping Tool. Ensure that you have the most current CCN that incorporates any change that have been made to your z13 order.

To enter the CCN and download the CFReport, complete the following steps:

1. Go to the following address:

   https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/chpidmain?Opendocument

2. Open the Tools window and, under Servers, select **CHPID Mapping Tool**.
3. Under Downloads, click **CFReport (I/O Configuration Data)**.
4. At the CFReport download page, enter your 8-digit CCN and click **Submit**.
5. Save the `nnnnnnnn.CFR` file to your workstation for later input into the CHPID Mapping Tool.

If you have any problems or require more information, click **Help about CFReport download**.

Availability mapping

When you are planning and configuring the z13 server, plan for maximum server and device availability in the event of single or multiple channel failures.

To help you configure a z13 to ensure maximum availability based on the characteristics of the server, the CMT has an Availability Mapping option. This option assigns channel paths to avoid a single point of failure. You also have the option of switching to manual mapping. In addition, you can map CHPIDs with the availability option and change paths manually afterward.

When you use the availability option, you must first provide a copy of the system’s IOCP source. Then, using the CHPID Mapping Tool, define priorities for the control units. The CMT can assign CHPIDs to the I/O ports and produce a new CHPID report that has maximum availability. This goal is achieved by distributing channel paths across different channel cards and different host channel adapters (HCAs and ICAs).

When you use the availability option within the CMT, you can assign priorities to control units. This process is useful when multiple control units provide a common function, such as multiple OSA control units that attach to the same network, or the control units for the primary and backup operator consoles. Priorities can also direct the CMT to pick PCHIDs for certain control units before others.

The CMT does not automatically map AIDs and ports to CIB or CS5 CHPIDs. You must complete this task manually. However, through the availability function, the CMT provides intersect information if multiple CIB or CS5 CHPIDs are defined to a control unit on the same HCA or ICA.

The PCHIDs, which are fixed on the z13, are mapped to CHPIDs by the CMT and subsequently assigned in the IOCP deck. The modified IOCP deck is an output of the CMT processing. It must be re-migrated into the validated work IODF that was used to generate the input IOCP deck used by the CMT. Other outputs of the CMT are tailored reports for your reference and for your IBM System Service Representative.

Requirement: The IOCP source file that is modified from the CHPID Mapping Tool must be migrated back into HCD. It cannot be used directly by IOCP. Any attempt to use the IOCP source that is created by the CHPID Mapping Tool in the IOCP program fails.

3.1.6 Other tools

The following tools might also be useful when you plan a new installation or MES:

- Batch Network Analyzer and Compression
- Worldwide Port Name (WWPN) Prediction Tool
- Coupling Facility Structure Sizer (CFSizer)
- Open System Adapter/Support Facility (OSA/SF)
- Processor Capacity Reference for IBM z Systems (zPCR)
- Power Estimation Tool
- Machine information

Batch Network Analyzer and Compression

IBM z Systems Batch Network Analyzer is a no charge, “as-is” tool that analyzes batch windows. It is available to customers, IBM Business Partners, and IBMers. This PC-based tool provides graphical and text reports, including Gantt charts and support for Alternate Processors.

By post-processing customer-provided SMF records, you can identify jobs and their BSAM/QSAM data sets, which are zEDC compression candidates across a specified 24-hour time window, typically a batch window. This helps you estimate utilization of a zEDC feature and the number of features required.

The tool is available from z Systems Washington Systems Center (WSC) CPS Tools team at the following websites:

- Customers:

  http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5132
- IBM Business Partners:

  https://www.ibm.com/partnerworld/wps/servlet/mem/ContentHandler/tech_PRS5133
- IBMers:

  http://w3.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS5126

Note: This tool replaces the BWATOOL.

Worldwide Port Name (WWPN) Prediction Tool

The Worldwide Port Name Prediction Tool for z Systems Fibre Channel Protocol (FCP) Channels helps prepare configuration files that are required or generated by the z Systems 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 worldwide port names (WWPNs), which uniquely identify physical or virtual Fibre Channel ports. They are typically used in Fibre Channel 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 which 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 at IBM Resource Link[1] and is applicable to all FICON channels defined as CHPID type FCP (for communication with SCSI devices) on z13. You can access the tool at the following web page:

https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/wwpnMain?OpenDocument&pathID=

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 PCHID values (for physical WWPNs). When a customer purchases a new “replacement” CPC and uses a “push-pull” (maintains current I/O Configuration definitions), the I/O serial number changes and most likely so do the PCHID values.

This means that the FCP WWPNs will change and the customer will need to reconfigure zoning in their storage area network (SAN) switches and LUN masking in their storage controllers. This is a a tedious, time-consuming, and error-prone process.

A solution was introduced with the z13. This feature in the order and fulfillment process enables a client to have the option to keep the same I/O serial number on their replacement CPC to facilitate WWPN persistence. This feature provides a means to export PCHID values of new system into `.csv` file, assign “alias” definitions to them (to match old system PCHID values) and import the edited file into the order fulfillment process.

[1] To access the IBM Resource Link, registration is required.

Coupling Facility Structure Sizer (CFSizer)

The Coupling Facility Structure Sizer Tool (CFSizer) is used to monitor various aspects of Coupling Facility partitions. It is useful for gathering information about structure sizes, usage, potential bottlenecks and coupling link availability.

Use the CFSizer to plan the amount of storage that must be allocated for coupling facility partitions more accurately.

Open System Adapter/Support Facility (OSA/SF)

The OSA/SF is an application that helps you customize and manage your OSA-Express features. It also allows you to obtain status and operational information about the HCD-defined OSA-Express features and helps with problem determination. OSA/SF is required for configuring the OSA-EXPRESS features on a OSE CHPID type, which functions in non-queued direct I/O (non-QDIO) mode. For more information about OSA/SF, see 14.1, “Configure OSE channel with Open System Facility (OSA/SF) on the HMC” on page 654.

Processor Capacity Reference for IBM z Systems (zPCR)

The zPCR is a PC-based productivity tool that runs under Windows. It provides capacity planning insight for IBM z Systems processors that run various workload environments under z/OS, z/VM, and Linux. Capacity results are based on IBM Large Systems Performance Reference (LSPR) data supporting all IBM z Systems platforms and LSPR data.

Power Estimation Tool

This tool estimates the power consumption for the specified configuration. The object of the tool is to produce an estimate of the power requirements to aid you in planning for your installation.

Machine information

Machine information is a set of reports that are based on data that the system sends to IBM as part of its IBM maintenance agreement.

Profiling is done against your Customer Number so be sure that you have it available before you register for machine information.

To access the machine information, your IBM ID must be authorized first. Do so by clicking **Register for machine information** on the Machine information web page:

https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/machineinformation?OpenDocument

Then, complete the following steps:

1. Under Servers in the Tools window, click **Machine information**.
2. Click **View all machines**.

If none of your systems are accessible, request access from your IBM Customer Representative.

If your IBM ID is registered with your Customer Name and Number, you get a list of servers that have been profiled for your Customer Number. Click the serial number of the server about which you are inquiring.

Reports that are available include information about client data, system status, engineering change/microcode change level (EC/MCL) and CHPIDs.

Further information about tools

Table 3-2 provides references to more detailed descriptions of the tools.

Table 3-2 IBM website addresses for the tools

Tool	Website for more information
HCD and HCM	http://www.ibm.com/systems/z/os/zos/features/hcm
CMT	https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/chpidmain?Opendocument
WWPN	https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/wwpnMain?OpenDocument*&pathID=
CFSizer	http://www.ibm.com/systems/z/cfsizer/
z/PCR	http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381
OSA/SF	http://www.ibm.com/systems/z/os/zos/zos_elefeat.html#o
Power Estimation Tool	https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/powerestimationmain?OpenDocument
WWPN Prediction Tool	https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/wwpnMain?OpenDocument&pathID=
Machine information	https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/machineinformation?OpenDocument

3.2 Hardware Management Console and Support Element configuration planning

The Hardware Management Console (HMC) communicates with each central processor complex (CPC) through the CPC's Support Element (SE). When tasks are run at the HMC, the commands are sent to one or more SEs. The SEs then issue commands to their CPCs. CPCs can be grouped at the HMC so that a single command can be passed along to all of the CPCs defined to the HMC. One HMC can control up to 100 SEs and one SE can be controlled by 32 HMCs.

An HMC can also communicate directly to the SE of a zBX model 004. In this instance, a pair of HMCs are required (primary and alternate) and are assigned an Ensemble identity. For more information, see 3.2.11, "Connectivity for the HMC and the SE with ensemble enabled" on page 91.

The SEs are now two integrated System X servers (x3550 M4 server) that are supplied with the z13. The SEs are closed systems, as are the HMCs, and no other applications can be installed on them.

The microcode for the z Systems platform and zBX is managed by the HMC.

The HMC Remote Support Facility (RSF) provides communication to a centralized IBM support network for hardware problem reporting and service. See 3.2.7, "Remote Service Facility (RSF)" on page 69.

3.2.1 Support Elements (SEs)

The SEs supplied with the z13 are two integrated 1U x86 servers. Both units are installed at the top of the A frame. One is a primary SE and the other is an alternate SE. The primary SE is the active one; the alternate acts as the backup. The SEs are closed systems so no other applications can be installed on them. See Figure 3-3.

As part of the z13, the SEs now consist of the following components:

- DC power supply
- Customized keyboard/display for service use only
- DVD drive
- USB attached Smart Card Reader always present to be used by Flash Express and Feature on Demand (FoD)
- Display/Keyboards: Two IBM 1U 18.5-inch Standard Consoles (1723-8BX) “Gate mounted”

Figure 3-3 Support Element - Display/Keyboard “gate”

3.2.2 Hardware Management Console Application V2.13.0

The z13 server includes version 2.13.0 of the HMC. Table 3-3 on page 52 shows the HMC version history.

Table 3-3 HMC and SE: z Systems HMC version history

Machine family	Machine type	Firmware driver	Version	Ensemble Node Potential
z13	2964	22	2.13.0	Yes
zBX Node	2458 Model 4	22	2.13.0	Required
zBC12	2828	15	2.12.1	Yes
zEC12	2827	15	2.12.1	Yes
z114	2818	93	2.11.1	Yes
z196	2817	86	2.11.0	Yes
z10 BC	2098	79	2.10.0	Yes
z10 EC	2097	79	2.10.0	No
z9 BC	2096	67	2.9.2	No
z9 EC	2094	67	2.9.2	No

HMC feature code (FC) changes

Available with the z13 are two types of HMCs: FC 0092 standard deskside HMC or 0094, which is the new rack-mounted HMC. Older HMCs (FC 0091 or FC 0092) can be carried forward to control a z13.

Order information:

- FC 0094 can be ordered only with the z13, but not older CPCs
- For FC 0094, the customer must supply the rack and appropriate AC power.
- Up to 10 HMCs are orderable.

As with the zEC12, an Ensemble requires a primary and alternate HMC pair to support z13. An identical pair is required (Two FC 0094, two FC 0092 or two FC 0091).

Notes:

- A service (orderable by IBM) is available to upgrade HMC FC 0092 or FC 0091 features of another z Systems platform to HMC Driver 22 LIC to support a z13. HMC application Driver 22 will support z196 (N-4) and later only.
- MES is required for existing HMCs with code at Driver 93 or lower to increase RAM from 8 GB to 16 GB.

New backup options are available with HMC 2.13. Critical z13 HMC data can be backed up using USB storage and using FTP/Secure FTP. Critical z13 Support Element (SE) data can be backed up to the SE, Alternate SE Hard Drive, or by using FTP or Secure FTP.

Older machine’s HMC and SE will have USB storage only. A new, optional 32 GB USB “stick” (FC 0848) is offered if needed to accommodate the increased data from Unified Resource Manager, zBX, IBM zAware and other growing code/data needs.

Note: As with the zEC12, 1000BASE-T switches are no longer offered. A preferred alternative is to use compatible, customer-provided 1000BASE-T switches.

3.2.3 HMC access

Starting with HMC Version 2.11.0, the Hardware Management Console application offers only one type of access called *full function remote access*. This single type of access provides the full complement of tasks and functions to all users.

If you need occasional monitoring and control of managed objects that are connected to a single local HMC, a web browser is a good choice. An example of using a web browser might be off-hours monitoring from home by an operator or system programmer.

Each HMC contains a web server that can be configured to allow remote access for a specified set of users. If a client firewall exists between the web browser and the local HMC, use the ports that are listed in Table 3-4. These ports allow the web browser to communicate with an HMC.

Table 3-4 HMC and SE: Port Numbers for HMC communication

Port	Usage
TCP 443	Secure browser to web server communication
TCP 9960	Browser applet communication
TCP 9950-9959	Proxy support for Single Object Operation

With an HTML-based user interface and a supported web browser, a remote user operates in the same manner as a local user who is working at the HMC itself.

Although the same user interface is used locally and remotely, the sets of tasks available to local and remote users are not identical. For example, the Format Media task is available only to a local user because it involves inserting the media into a physical HMC drive. This task is not available to a remote user in tree style. However, if you are using the classic style, an error results, as shown in Figure 3-4.

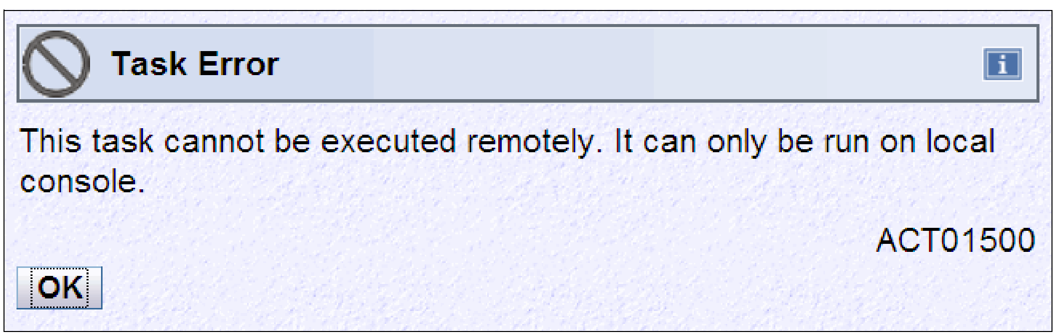

Figure 3-4 HMC and SE: Task Error message

Customize console services

Global enablement of remote access is controlled through the Customize Console Services task, as shown in Figure 3-5 on page 54. However, individual user level remote access enablement is managed using the user profiles console actions task that is shown in Figure 3-6 on page 55.

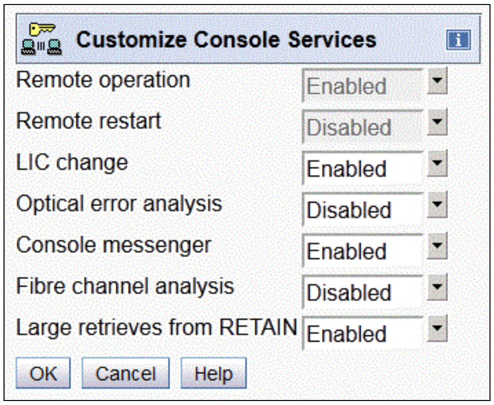

Figure 3-5 HMC and SE: Customize Console Services

The console services can be customized as follows:

- Remote operation

  Use this service to control whether this HMC can be operated by using a web browser from a remote workstation.

- Remote restart

  Use this service to control whether this HMC can be restarted by a user who accesses it from a remote workstation. If this service is Disabled, only local users at this HMC can use the Shutdown or Restart task.

- LIC change

  Use this service to control whether this HMC provides change management operations for its defined objects and for other HMCs. This option must be enabled for an HMC connected to both the local and remote processors for HCD to be able to write an IOCDS to a processor on which it is not running.

- Optical error analysis

  Use this service to control whether this HMC analyzes and reports optical problems for its defined objects. Optical problems are problems that occur on coupling facility channel links.

- Console messenger

  Use this service to control whether the console messenger facility is active on this HMC or not. The console messenger facility allows users of this HMC to send and receive instant messages. They can also broadcast messages to other users of this console and remote consoles.

- Fibre Channel analysis

  Use this service to control whether this HMC analyzes and reports Fibre Channel problems.

- Large retrieves from IBM RETAIN®

  Use this service to control whether this HMC can retrieve internal code changes from RETAIN for Engineering Change (EC) streams that are expected to contain a large amount of data.

- Check held MCLs during install

  Controls whether this console checks RETAIN for any MCLs on hold when an installation and activation is run. Use Enabled to prevent activation of released fixes that have been discovered to have problems.

Note: The tasks that require removable media cannot be run remotely. Also, Remote Operation and Remote Restart options are not available if you access the HMC remotely.

Customize user profiles

Lightweight Directory Access Protocol (LDAP) support for HMC user authentication allows an HMC to be configured to use an LDAP server to run user ID and password authentication at logon. The user ID is defined on the HMC along with the roles to be given to the user ID. HMC settings that are related to the user ID are on the HMC, and the LDAP directory is used to authenticate the user. This configuration eliminates the need to store the user ID's password locally. Both Secure Sockets Layer (SSL) and non-SSL connections to the LDAP server are supported.

The LDAP configuration window is shown in Figure 3-6.

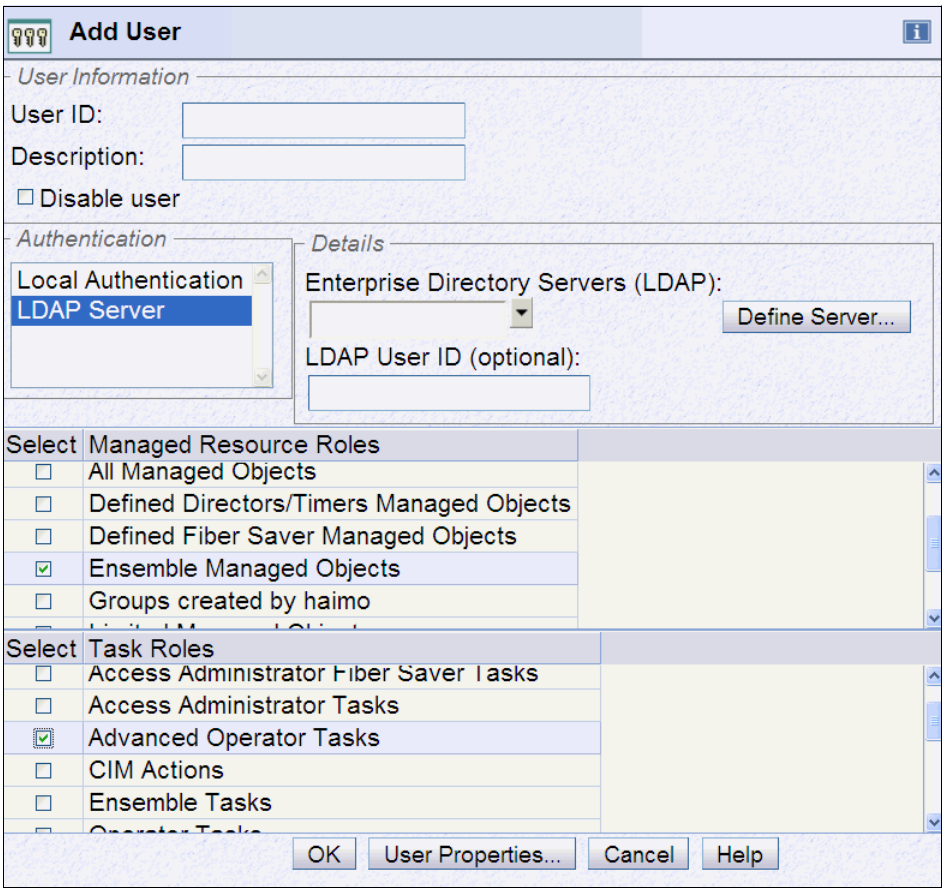

Figure 3-6 HMC and SE: Add User window

Individual user remote access enablement is accomplished by selecting **Allow remote access via the web** on the User Properties window (Figure 3-7). If you want to use the Web Services API, **Allow access to management interfaces** must be selected. For a new user profile, both check boxes are clear, by default.

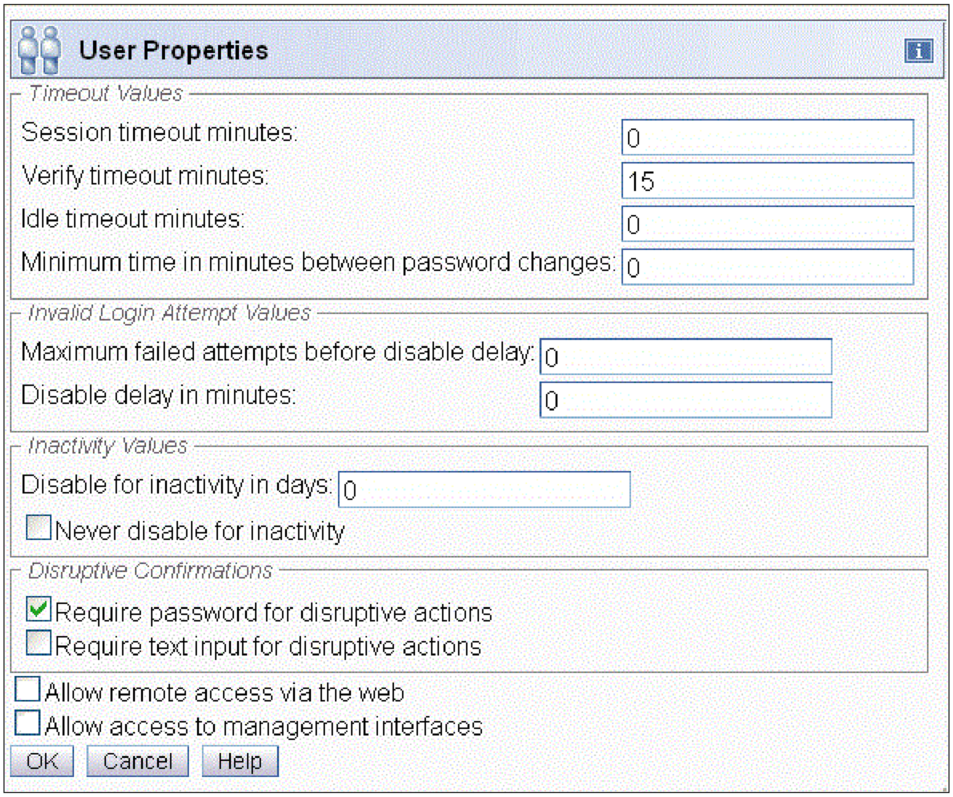

Figure 3-7 HMC and SE: User Properties window

3.2.4 Application enhancements

The z13 CPC includes the new HMC application Version 2.13.0. Use the “What’s New Wizard” to explore new features available for each release. See Figure 3-17 on page 64.

The application provides enhancements to the following areas:

- The SE, HMC, and TKE hardware
- Alternative to USB flash drive
- uEFI/BIOS enhancements
- Additional RSF infrastructure changes

System Control Hub

The System Control Hub (SCH) replaces the Bulk Power Hub. It supports 1Gbps and has redundant physical networks. The SE is always connected to the System Control Hub. Customer supplied Ethernet switches can be connected to the System Control Hubs (two switches recommended). Other server’s SEs and HMCs may be connected to the Ethernet switches.

Alternative to the USB Flash Drive

For z13 and zBX model 004, FTP and Secure FTP can be used instead of using a USB flash drive. The FTP server must be supplied by the customer. If used, it must be set up in advance of the backup. A Backup FTP Site Access Information (Figure 3-8) task is added to the HMC.

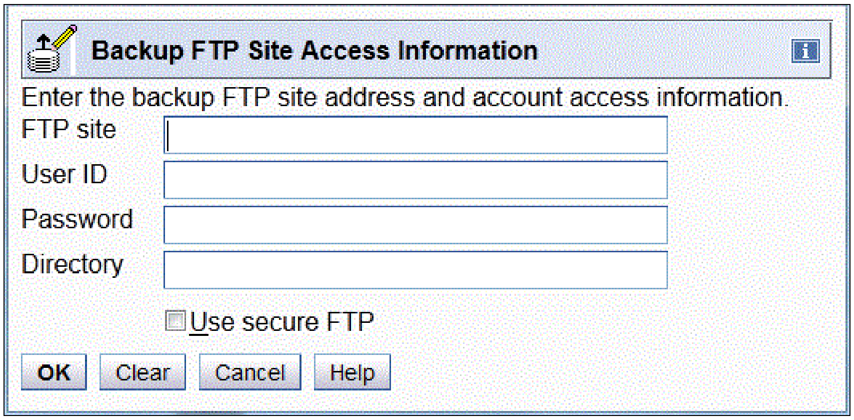

Figure 3-8 HMC Configure Backup Settings

When backing up the SE Critical Data on the z13 and zBX Model 004, there is no USB option. The data must be saved to the Primary SE HDD and Alternate SE HDD or to the FTP Server. See Figure 3-9.

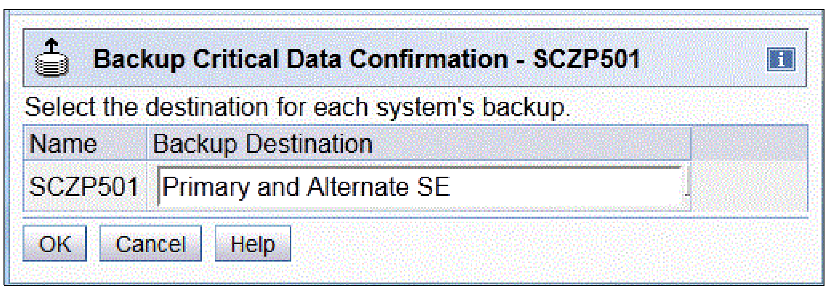

Figure 3-9 HMC backing up the SE critical data

When backing up console data from the HMC, three options are available (Figure 3-10):

- USB
- FTP server
- USB and FTB server

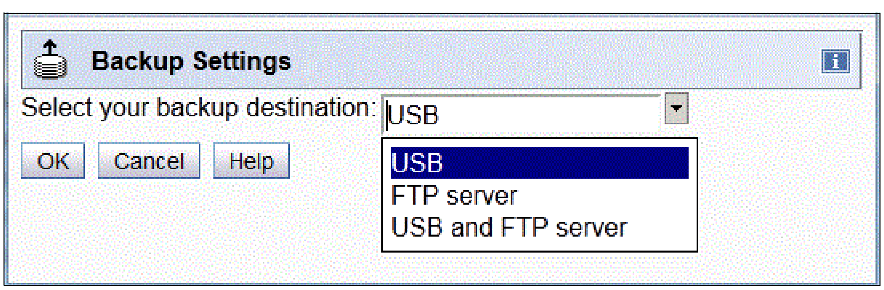

Figure 3-10 HMC backing up the Critical Console Data

For older CPCs (zEC12, z196 and older), critical data must be saved only to USB. If multiple CPCs are being backed up, a large 32 GB USB (FC0848) might be required because the growth and addition of new components such as zBX, IBM zAware, and others.

All options shown (Figure 3-11) can be a scheduled operation if you want.

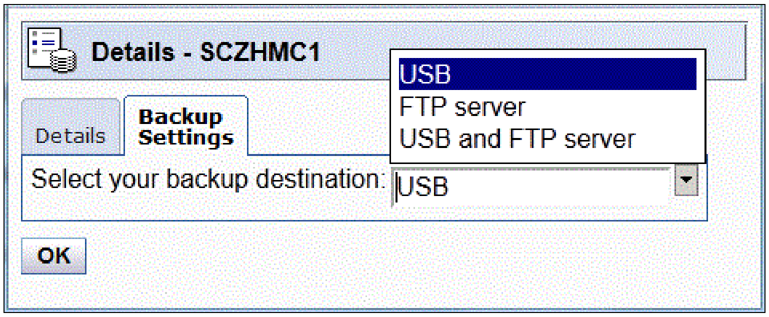

Figure 3-11 Customize scheduled operations

Firmware: uEFI/BIOS change

With HMC version 2.13, you can disable the ability to boot from removable media. You might want to protect against unauthorized booting of an OS on bootable removable media. An admin password can be set for the uEFI/BIOS. This can be accomplished by making a uEFI/BIOS configuration change by an IBM System Service Representative (SSR). If the admin password was set, the SSR will need customer input for the service actions where boot from media needs to be enabled.

Additional RSF infrastructure changes

HMC 2.13.0 will continue the enhanced IBM Support System for Remote Support Facility (RSF) call-home. HMC 2.13.0 will support the same functions as in HMC 2.12.1. In addition, the following new functions are supported:

- e-business on demand (eBoD) to support all Capacity on Demand records that can be ordered from Resource Link (CIU Permanent, On/Off CoD, CPE, CBU)
- Fix delivery with full fix deliverable capabilities
- Problem Management Viewable (PMV) “update” support

For more information about RSF, see 3.2.7, “Remote Service Facility (RSF)” on page 69

Note: Driver 22 on the HMC and SE is required for use of any of enhancement to RSF.

3.2.5 Configuration and services

This section summarizes several enhancements to the HMC and SE panels to support Flash Express, IBM zAware, Server Time Protocol (STP), and other HMC topics.

Flash Express

Flash Express introduces solid-state drive (SSD) technology to the z Systems family. An operating system, such as z/OS, can access blocks of flash memory as storage locations within a logical partition. For more information about Flash Express, see 13.1, “Overview of Flash Express” on page 632

IBM z Advanced Workload Analysis Reporter (IBM zAware)

IBM zAware is an integrated, self-learning, analytics solution for IBM z/OS that helps identify unusual system behavior in near real time. It is designed to help improve problem determination for support personnel so they can restore service quickly and improve overall availability.

For more information about the configuration and setup of the IBM zAware management function, see either of these publications:

- *Extending z/OS System Management Functions with IBM zAware*, SG24-8070
- *System z Advanced Workload Analysis Reporter (IBM zAware) Guide,* SC27-2623

Server Time Protocol (STP) NTP support with broadband security

HMC level 2.13.0 continues to support the possibility to use HMC NTP authentication. The SE NTP support is unchanged. To use this option on the SE, configure the HMC as an NTP server for the SE.

The Network Time Protocol (NTP) client support allows an STP-only Coordinated Timing Network (CTN) to use an NTP server as an External Time Source (ETS).

Exception: The ETS connection through modem is not supported on the z13 HMC.

Authentication support with a proxy

Some client configurations use a proxy to access outside the corporate data center. NTP requests are User Datagram Protocol (UDP) socket packets and cannot pass through the proxy. The proxy must be configured as an NTP server to get to target servers on the web. Authentication can be set up on a proxy to communicate to the target time sources.

Authentication support with a firewall

Some clients use a firewall. HMC NTP requests can pass through the firewall. If you use this configuration, use HMC authentication to ensure untampered time stamps.

Symmetric key and autokey authentication

With the symmetric key and autokey authentication, the highest level of NTP security is available. HMC level 2.13.0 provides windows that accept and generate key information to be configured into the HMC NTP configuration. You can also issue NTP commands as shown in Figure 3-12 on page 60.

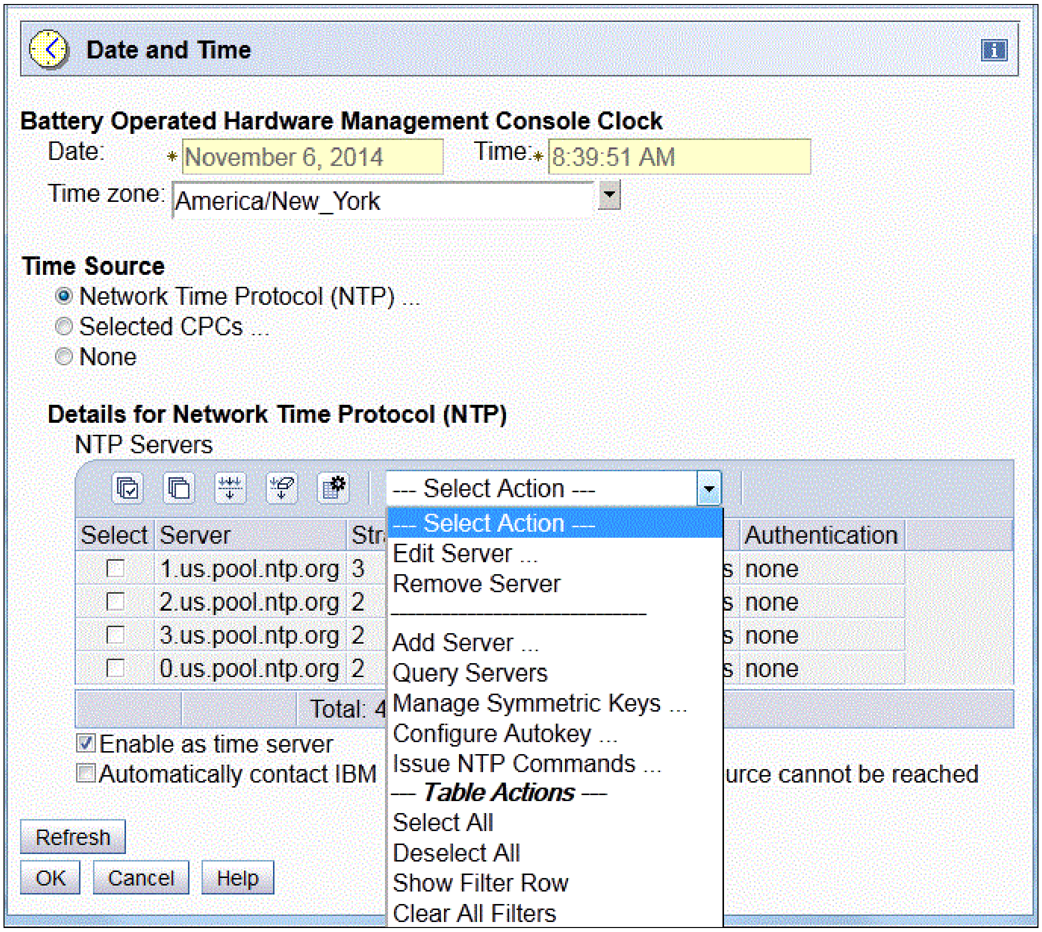

Figure 3-12 HMC NTP broadband authentication support

Symmetric key and autokey authentication option are as follows:

- Symmetric key (NTP V3-V4) authentication

  Symmetric key authentication is as described in RFC-1305 (made available in NTP Version 3). Symmetric key encryption uses same key for both encryption and decryption. Users exchanging data keep this key to themselves. Messages encrypted with a secret key can be only decrypted with the same secret key. Symmetric does support network address translation (NAT).

- Symmetric key autokey (NTP V4) authentication

  Autokey uses public key cryptography as described in RFC-5906 (made available in NTP Version 4). The key generation for the HMC NTP is done by clicking **Generate Local Host Key** in the Autokey Configuration window. Doing so issues the **`ntp-keygen`** command to generate the specific key and certificate for this system. Autokey authentication is not available with NAT firewall.

- Issue NTP commands

  NTP command support is also added to display the status of remote NTP servers and the current NTP server (HMC).

For more planning and setup information for STP and NTP, see Chapter 9, “Server Time Protocol (STP)” on page 511. In addition, 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

Time coordination for zBX components

With zBX Model 004, separate SEs control CPC and zBX components independently. SEs managing zBX has no direct tie to any CPC. zBX Model 004 will synchronize time with a primary NTP-configured HMC or selected CPCs as its time source. zBX SE automatically retrieves the time from the HMC and then synchronizes the time to the zBX blades and BladeCenters. This configuration is shown in Figure 3-13.

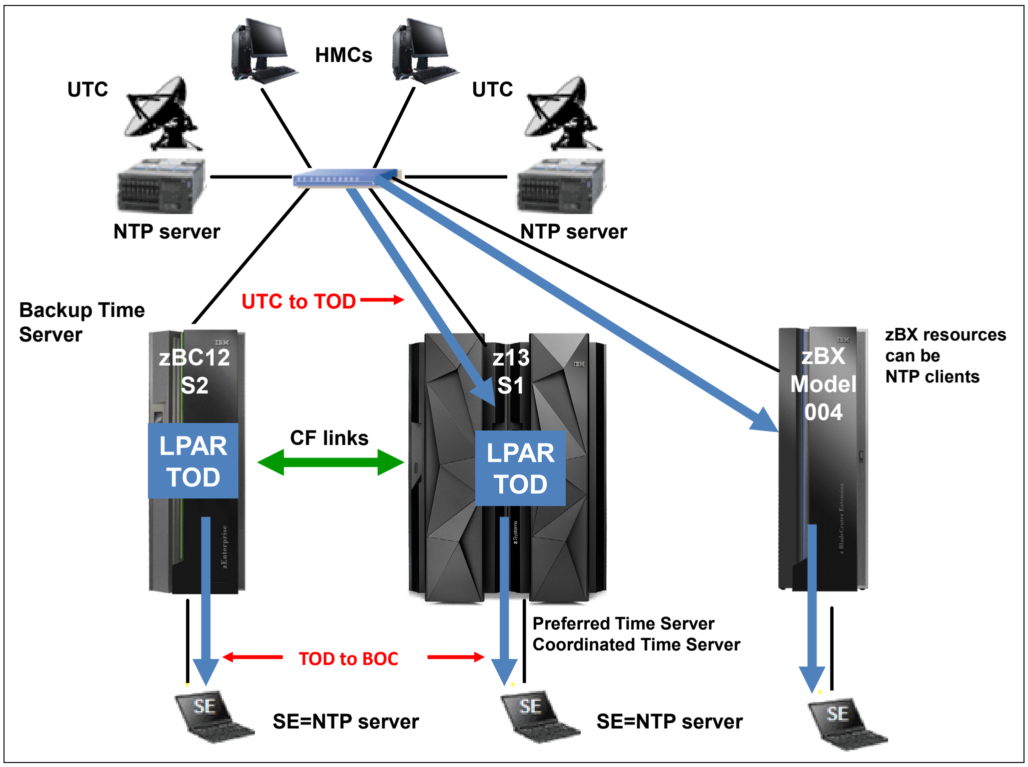

Figure 3-13 Time coordination for zBX model 004 +components

Microcode installation by MCL bundle target

A bundle is a set of MCLs grouped during testing and released as a group on the same date. z13 enhances the support to allow installation and activation to a specific bundle level, as shown in Figure 3-14.

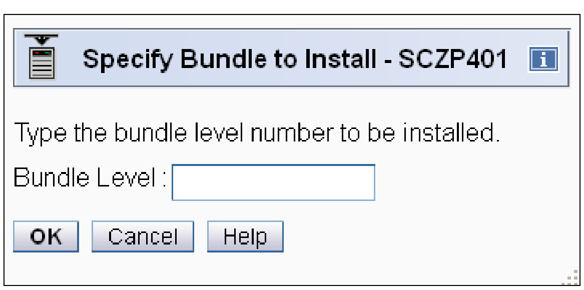

Figure 3-14 Specify Bundle to Install

The System Information window is enhanced to additionally show a summary bundle level for the activated levels. To do so, all applicable Engineering Change (EC) streams information must have the same bundle level. If they do not match, Not Available is displayed as the bundle level. An example of the bundle level is shown in Figure 3-15.

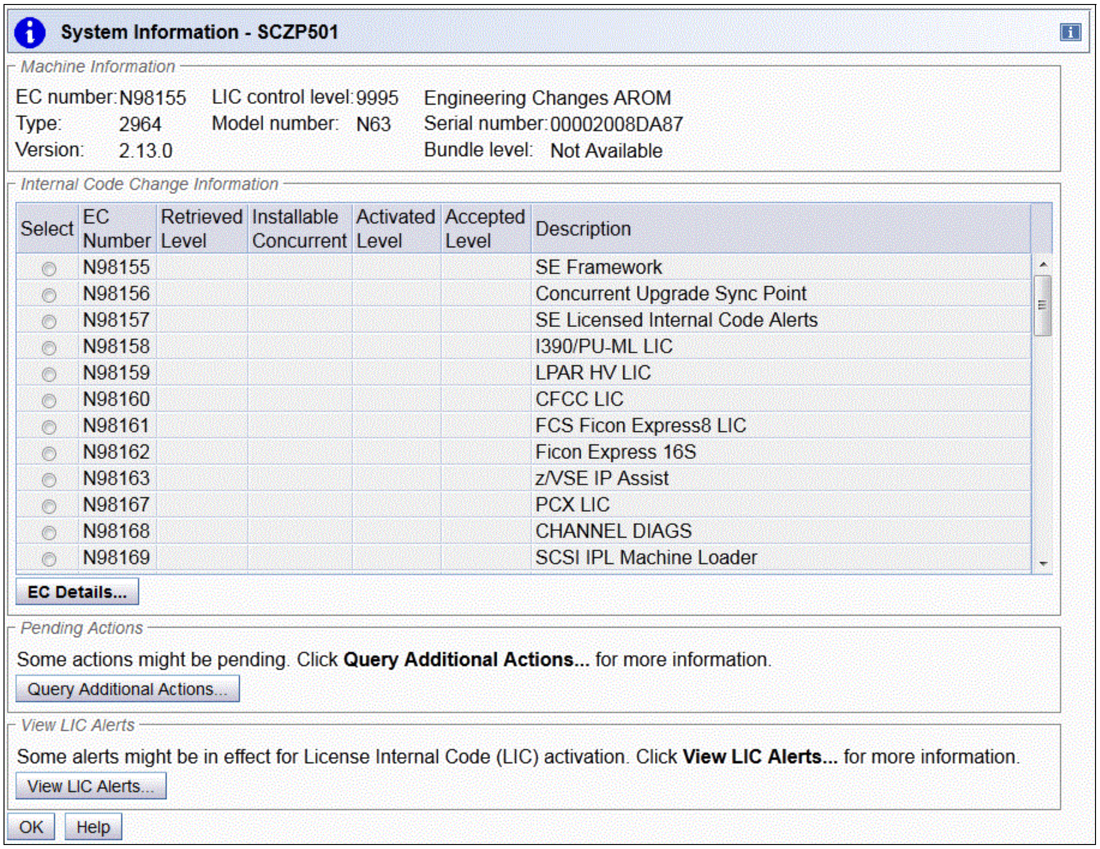

Figure 3-15 System information with bundle level

Removal of modem support from the HMC

This change impacts clients that once used modems for RSF or STP NTP access. For more information, see 3.2.7, “Remote Service Facility (RSF)” on page 69 and “Server Time Protocol (STP) NTP support with broadband security” on page 59.

Remove support for the Coprocessor Group Counter Sets

The System z10 GA2 server introduced various counters for use with the CPU Measurement Facility and Hardware Instrumentation Services. These counters are used to count events for each central processor (CP) such as cycles used and instructions. The Coprocessor Group Counter Set was the counter for crypto coprocessors that were shared by CPs. In the z13, each physical processor has its own crypto coprocessor and no longer must share this coprocessor with another CP. This removes the need for the Coprocessor Group Counter Set. All necessary crypto counter information is now available in the crypto activity counter set.

The check box for the Coprocessor Group Counter Sets was removed from the Image profile definition and the Change Logical Partition Security task, as shown in Figure 3-16.

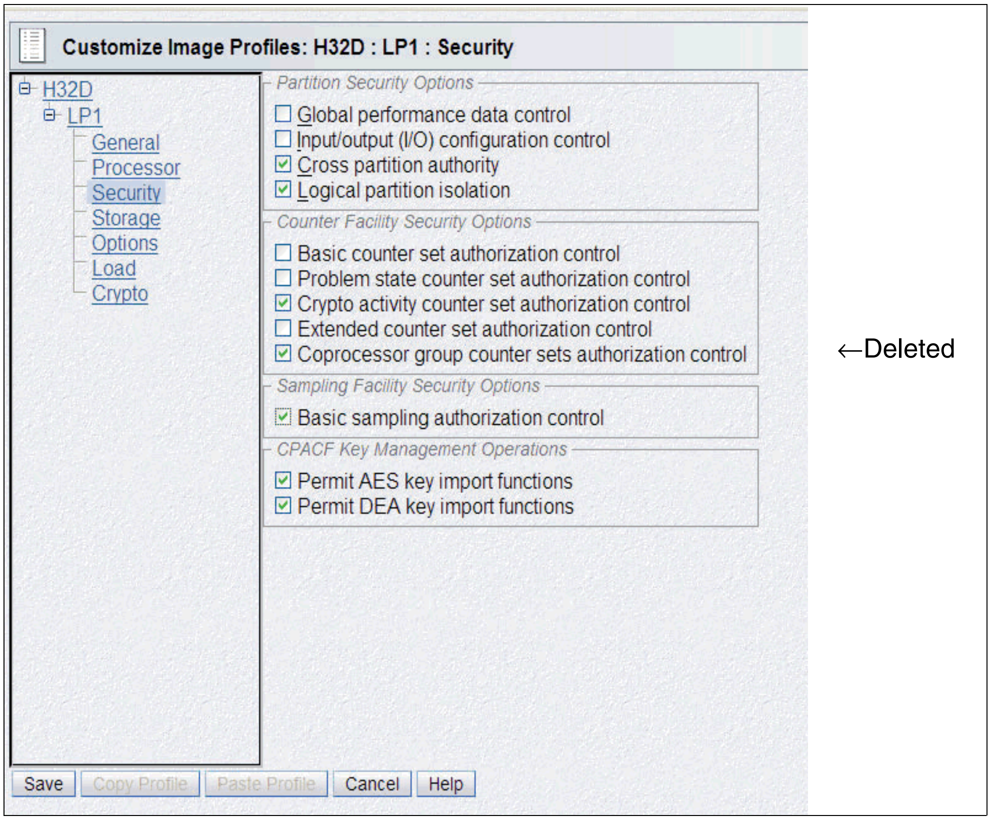

Figure 3-16 Remove support for the Coprocessor Group Counter Sets

User interface

Miscellaneous user interface enhancements were made in the z13:

- A simple wizard (named What's New) describes new features available on the HMC (Figure 3-17).

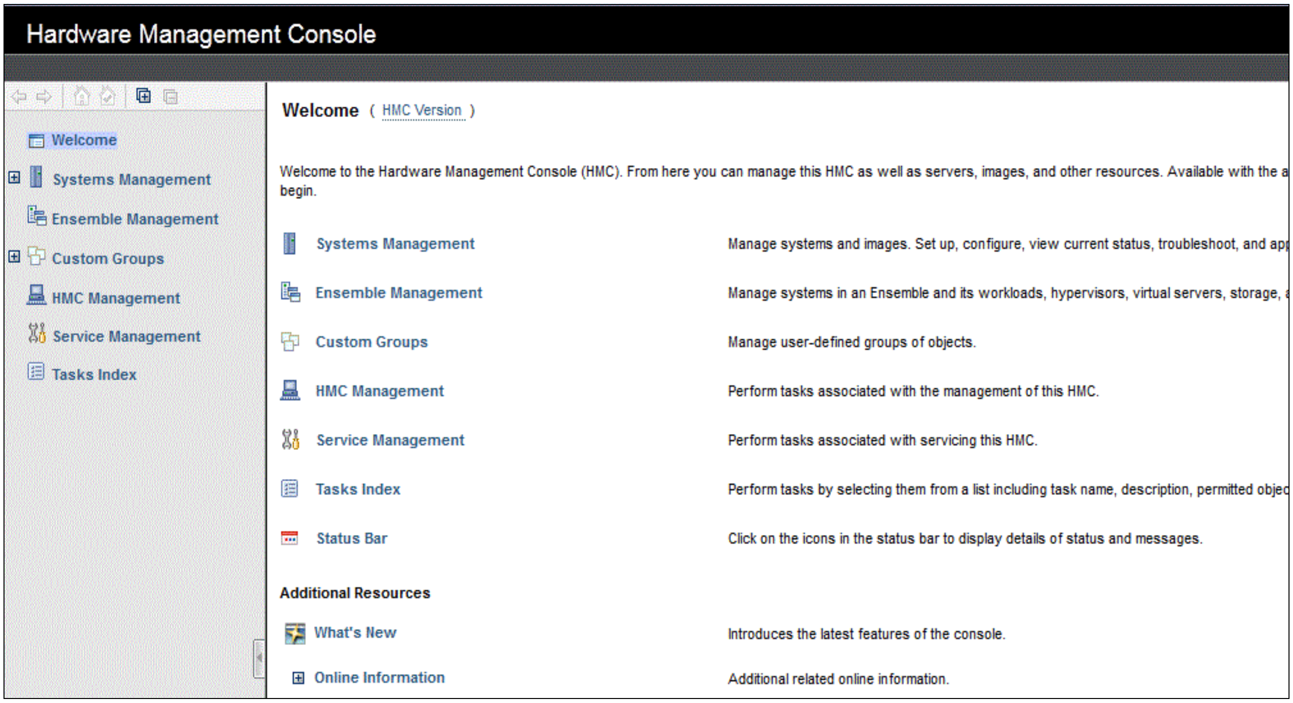

Figure 3-17 HMC and SE: What's New task

- Both the Tree Style User Interface and Classic Style User Interface are improved.
- Firefox Version 10 or newer is supported for accessing the HMC remotely. The following browsers are supported:
  - Internet Explorer Version 9 or newer
  - Firefox Version 10 or newer
  - Chrome Version 11 or newer

3.2.6 Application programming interface (API) customization

The purpose of the console APIs is to provide an open set of interfaces and a workstation platform for system management application providers. You can use the interfaces to use object-based industry-standard programming interfaces to collect the hardware information that you need to provide an integrated hardware and software system management solution.

Figure 3-18 shows the integration of system management applications by using the Console application open programming interfaces to provide a single system image (SSI) and a single point of control (SPOC).

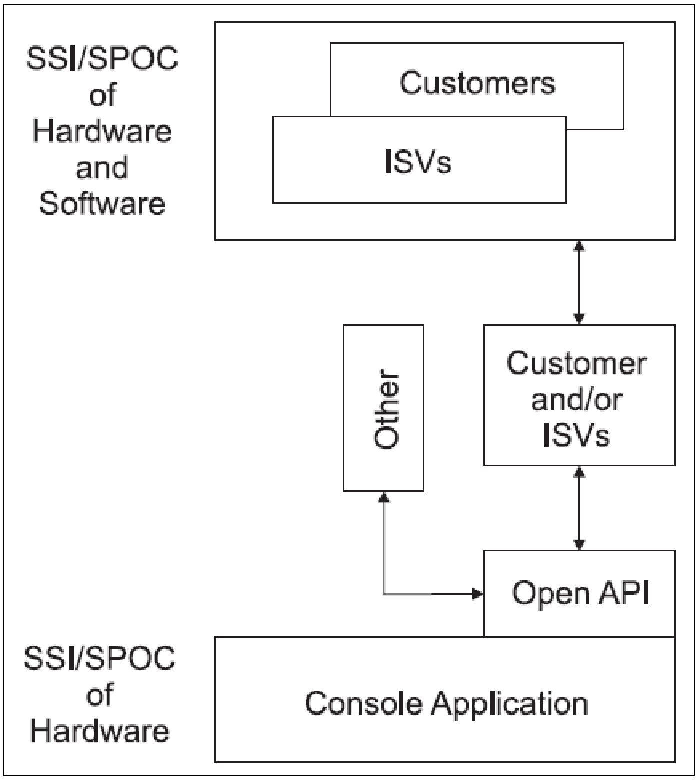

Figure 3-18 Console API's objectives

Simple Network Management Protocol (SNMP)

SNMP describes management data and the protocols for exchanging that data between heterogeneous systems. The protocols include both the description of the management data that are defined in the Management Information Base (MIB), and the operations for exchanging or changing that information. By implementing common protocols, management data can be exchanged between different platforms with relative ease.

SNMP defines an architecture that consists of these items:

- Network management applications
- Network management agents and subagents
- Network elements, such as hosts and gateways

The SNMP application can ask agents for specific information about network elements. Conversely, agents can tell the network management application when something happens to one or more network elements.

An SNMP agent can be configured for an HMC to allow other system management applications to manage the HMC applications by using HMC APIs. You must enable the SNMP APIs to support z13 features such as the z/OS Capacity Provisioning Manager (CPM) or the Active Energy Manager (AEM) that monitors the power and thermal data of a z13.

SNMPv3 is the new industry standard, and provides strong security mechanisms for authentication and privacy of communications. The new SNMP API client libraries now support up to SNMPv3 so that they include these strong security features.

Common Information Model (CIM)

The HMC allows external clients a limited ability to remotely manage system objects by using the CIM management architecture. CIM is defined and controlled by the Distributed Management Task Force (DMTF), which is an industry consortium. The DMTF publishes many white papers and specification documents that describe the many aspects of CIM. These documents, which include several that are referenced later, can be accessed at the DMTF website:

http://www.dmtf.org

Web Service API

The Unified Resource Manager is a collection of advanced hardware and virtualization management functions that are delivered as z Systems firmware. The functions of Unified Resource Manager are implemented as a cooperating set of components. These components are hosted on the HMC, the SE, the blades of a zBX, and as extensions to z/VM. It provides a uniform, integrated, and workload-oriented administrative model for the heterogeneous computing configuration that is provided by a z13 system. The Unified Resource Manager provides these functions:

- Hardware inventory, initialization, configuration, monitoring, and problem analysis for both a z Systems CPC (including both traditional z Systems computing resources) and a zBX.
- Firmware installation and update for the HMC, SE, traditional CPC components, zBX infrastructure, intraensemble data network (IEDN) elements, and zBX blades.
- Operational control and energy management for these hardware elements.
- Configuration of function-specialized workload accelerators available as blade optimizers in the zBX.
- Provisioning, configuration, control, and monitoring of virtualized computing systems (virtual servers) on the firmware-managed IBM blade and z/VM environments.
- Secure management of the IEDN through the provisioning and control of IEDN virtual networks.
- Automatic and workload-oriented performance optimization of the heterogeneous, virtualized environment.

The HMC serves as the administrative access point for Unified Resource Manager. In that capacity, the HMC provides a web-based remote graphical UI to make the Unified Resource Manager functions available to users. In addition, it hosts the implementation of Unified Resource Manager Web Service API (Web Services API).

The Web Services API is a web-oriented programming interface that makes the underlying Unified Resource Manager capabilities available for use through higher level management applications, system automation functions, or custom scripting. The functions that are exposed through the API support several important usage scenarios in virtualization management. These scenarios include resource inventory, provisioning, monitoring, automation, and workload-based optimization, among others.

To configure the API settings, you must log on to the HMC using the *Access Administrator* role, and select **HMC Management** → **Customize API Settings**.

Figure 3-19 on page 67 shows the Customize API Settings window, which is used to add SNMP information.

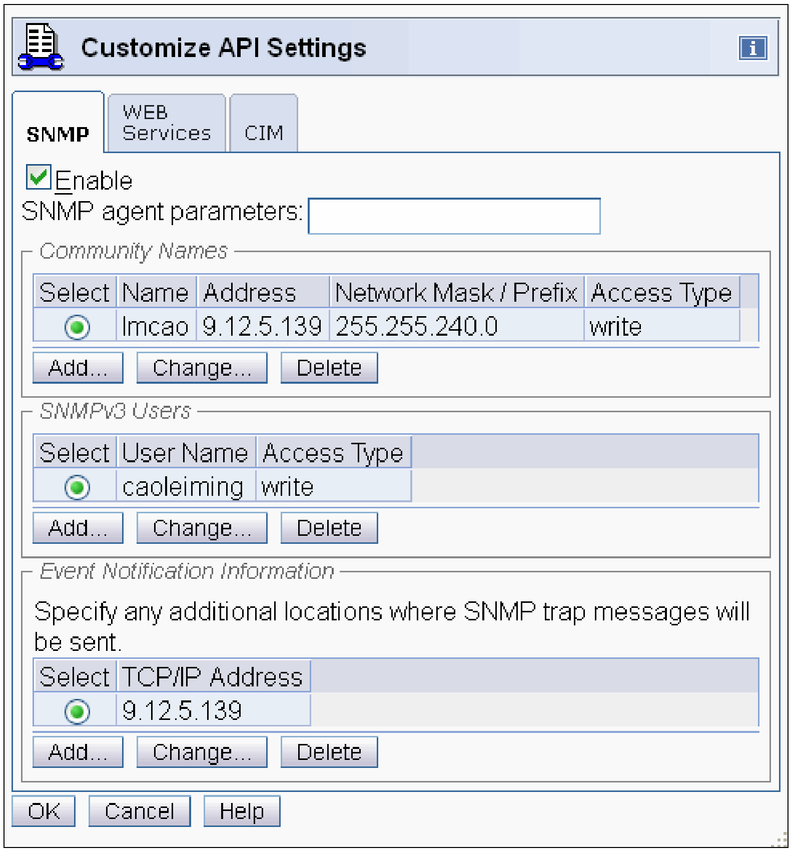

Figure 3-19 HMC and SE: Customize API Settings - SNMP

This access allows applications that were not supplied as part of the Hardware Management Console Application (HWMCA) to communicate with the objects defined to this HMC. You can use this task to enable or disable an SNMP agent. You can also set up a community name file and event notification information for an SNMP agent from the SNMP tab. You can enable or disable the Web Services API from the WEB Services tab. You can also enable or disable the CIM interface from the CIM tab.

To customize API settings, complete these steps:

1. From the SNMP tab you can enable SNMP APIs and add, change, or delete community names, SNMPv3 users and event notification information. See Figure 3-19.

   To allow applications to request information from the HMC, a community name must be added. Click **Add** under the Community Names section to specify one or more community names to use.

   Figure 3-20 on page 68 shows the Community Name Information window that is displayed after you click **Add**. The following information is required:

   Name Specifies the community name that is used to verify that a request for SNMP information is valid when a manager makes an SNMP request.

   Address Specifies the IPv4 or IPv6 IP address.

   Network mask Specifies a network mask that is logically ANDed with the IP address of the manager making an SNMP request.

   Access Type Specifies the access (read or write) you want to allow for SNMP requests.

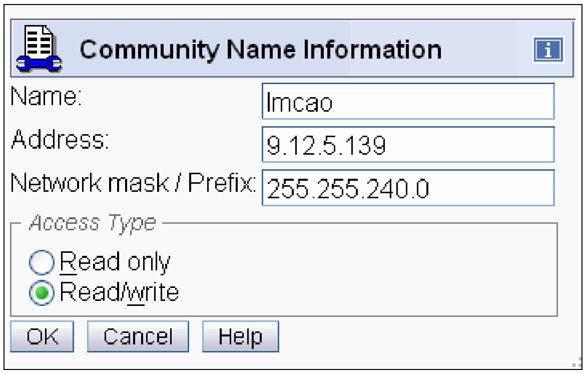

Figure 3-20 HMC and SE: Community Name Information

Community names and access type are specified by the application. For example, z/OS Capacity Provisioning Manager requires a defined community name and read/write access. System Director's AEM requires a defined community name, but read access is sufficient.

To allow applications to use SNMP APIs, click **Enable SNMP APIs**. After the SNMP APIs are enabled, the HMC must be restarted.

Specify any additional locations where enterprise-specific SNMP trap messages created by the Console should be sent in the **Event notification information** box. Entries can be added, changed, and deleted throughout by clicking **Add**, **Change**, and **Delete**. Adding entries to the **Event notification information** box causes the Console to send the specified event notifications to TCP/IP port 162 at the locations specified.

2. From the WEB Services tab, you can enable or disable the Web Services API and control the IP addresses and user access as shown in Figure 3-21.

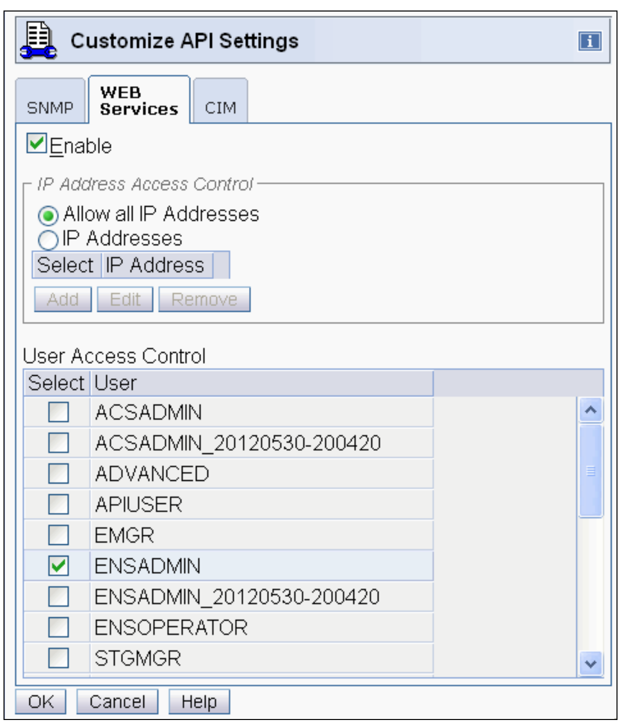

Figure 3-21 Customize API Settings: WEB Services

3. From the CIM tab, you can enable or disable the CIM interface as shown in Figure 3-22.

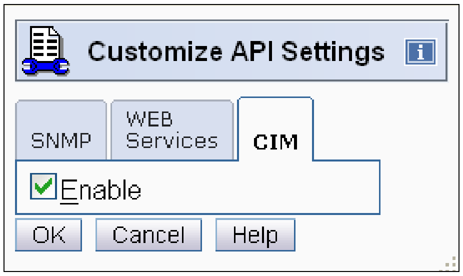

Figure 3-22 Customer API Settings: CIM

4. Click **OK** to save the SNMP, Web Services, or CIM configurations and continue.
5. Restart HMC.

Consideration: HMC is not restarted automatically. It must be done manually.

6. Configure the HMC user. In addition to the overall enablement on/off control and the optional client network address filtering, access to the API is further secured by the requirement for per-user authorization. By default, remote connections are not allowed for an HMC user. To enable remote connections, use the **User Profiles** or **Manage Users Wizard** task.

Note: By default, these tasks are available only to the ACSADMIN user.

7. Select the user for the CIM or Web Services clients, then select **Modify** from the User list. From the Modify User window, click **User Properties** as shown in Figure 3-7 on page 56.

   If you want to use CIM, select **Allow remote access via the web.**

   If you want to use the Web Services API, select **Allow access to management interfaces**.

   Click **OK**. Then, click **OK** on the main Modify User window.

   You can repeat the process for any other users that will be used by CIM or Web Services API clients.

   For more information about SNMP, Web Services, and CIM APIs, see these publications:

   – *System z Application Programming Interfaces*, SB10-7030
   – *Hardware Management Console Web Service API,* SB27-2616
   – *Common Information Model (CIM) Management Interfaces,* SB10-7154

3.2.7 Remote Service Facility (RSF)

The Hardware Management Console RSF provides communication to a centralized IBM support network for hardware problem reporting and service. The RSF provides these types of communication:

- Problem reporting and repairing of data
- Delivering fixes to the service processor and Hardware Management Console
- Hardware inventory data
- On demand enablement (optional)

You can configure the HMC to send hardware service-related information to IBM by using an Internet connection. Using an Internet connection provides these advantages:

- Significantly faster transmission speed
- Ability to send more data on an initial problem request, potentially resulting in more rapid problem resolution
- Reduced client expense (for example, the cost of a dedicated analog telephone line)
- Greater reliability

Unless your enterprise's security policy prohibits any connectivity from the HMC over the Internet, use an Internet connection.

The following security characteristics are in effect:

- Remote Support Facility requests are always initiated from the HMC to IBM. An inbound connection is never initiated from the IBM Service Support System.
- All data that is transferred between the HMC and the IBM Service Support System is encrypted in a high-grade Secure Sockets Layer (SSL) encryption.
- When it initializes the SSL encrypted connection, the HMC validates the trusted host using its digital signature that was issued for the IBM Service Support System.
- Data sent to the IBM Service Support System consists solely of hardware problems and configuration data. No application or client data is transmitted to IBM.

Internet connectivity

When an Internet connection is used, the HMC can be configured to use a second network card to physically separate a private LAN connection from the Internet-enabled network. The HMC can be enabled to connect directly to the Internet (Figure 3-23 on page 71) or to connect indirectly from a client-provided proxy server (Figure 3-24 on page 72). The decision about which of these approaches to use for your installation depends on the security and networking requirements of your enterprise.

Note: RSF through a modem is *not supported* on the z13 HMC. Modems that are on installed HMC FC 0091 hardware do not work with the HMC version 2.13.0, which is required to support z13.

All communications are handled through TCP sockets that are initiated by the HMC, and use a high-grade SSL to encrypt the data that is transmitted. The destination TCP/IP addresses are listed.

If using IPv4, firewall outbound connectivity must permitted to port 443 to the following destinations:

- 129.42.26.224 (traditional)
- 129.42.34.224 (traditional)
- 129.42.42.224 (traditional)
- 129.42.56.189 (enhanced)
- 129.42.58.189 (enhanced)
- 129.42.60.189 (enhanced)

If using IPv6, firewall outbound connectivity must permitted to port 443 to the following destinations:

- 2620:0:6c0:1::1000
- 2620:0:6c1:1::1000
- 2620:0:6c2:1::1000
- 2620:0:6c0:200:129.42.56.189 (enhanced)
- 2620:0:6c1:200:129.42.58.189 (enhanced)
- 2620:0:6c2:200:129.42.60.189 (enhanced)

If using an SSL Proxy and the plan for it is to resolve host names, it must accept the following host names:

- `www-945.ibm.com` (traditional)
- `esupport.ibm.com` (enhanced)

Remember: HTTPS port 443 is used for all communications.

If your HMC can be connected to the Internet and the external firewall can be set up to allow established TCP packets to flow outbound to the destination addresses, you can use a direct Internet connection. The use of Source Network Address Translation (SNAT) and masquerading rules to mask the HMC's source IP address are both acceptable (Figure 3-23).

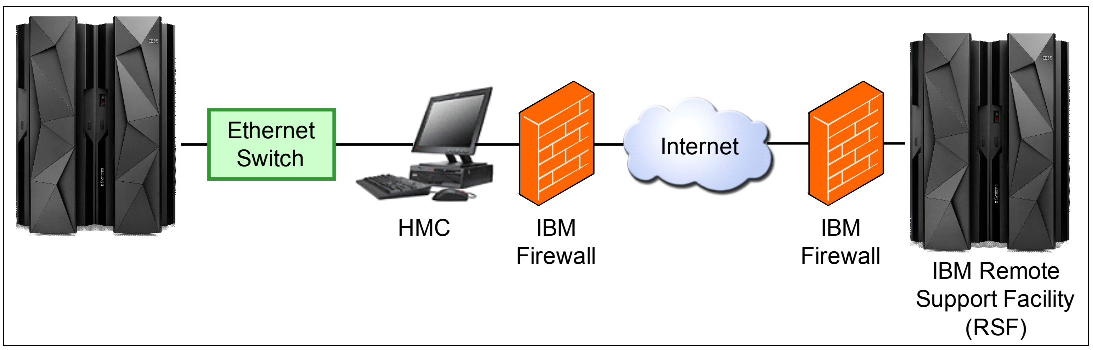

Figure 3-23 HMC and SE: Direct Internet connection

If your installation requires the HMC to be in a private network, you might be able to use an indirect Internet connection by using an SSL proxy. The proxy forwards requests to the Internet. One of the other potential advantages of using an SSL proxy is that the proxy can support logging and audit facilities. To forward SSL sockets, the proxy server must support the basic proxy header functions (as described in RFC 2616) and the CONNECT method. Optionally, basic proxy authentication (RFC 2617) can be configured so that the HMC authenticates before attempting to forward sockets through the proxy server as shown in Figure 3-24 on page 72.

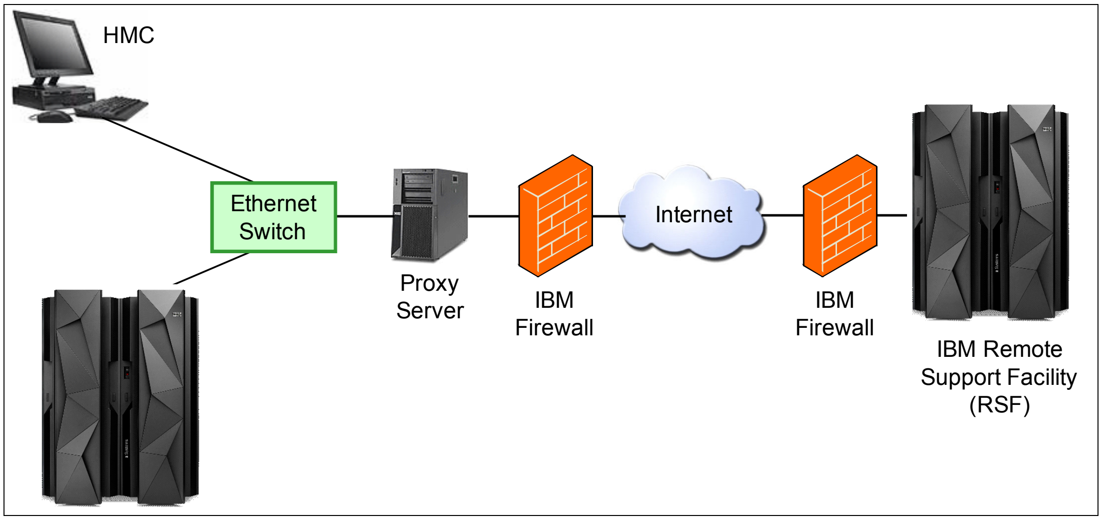

Figure 3-24 HMC and SE: Indirect Internet connection with a proxy server

Configuring outbound connections

At a high level, the following procedures are used to configure outbound connections:

1. Customize client information.
2. Customize remote services.
3. Customize outbound connectivity.
4. Customize network settings.
5. Customize console services.

Details for these procedures are provided in the following sections.

> **Requirement:** All steps that are needed to configure outbound connectivity require client approval.

Configuring an RSF connection

These procedures enable the HMC to connect to the RSF with a direct Internet connection. Begin at the HMC Settings window and complete the following steps:

1. Select **Customize Customer Information** and complete these steps:
   a. Enter the Administrator information.
   b. Enter System information.
   c. Enter Account information.
   d. Close the window.

   Completing these entry fields for each managed system allows your service structure to record necessary contact information.
2. Select **Customize Remote Services**.

   Remote service is two-way communication between the console and the IBM Service Support System (commonly known as RETAIN) to conduct automated service operations. Using remote service reduces the operator interaction that is needed to complete some service operations. It also provides some console tasks with another source or destination for sending or receiving service information.

In the window (Figure 3-25), select these options and then click **OK**.

a. **Enable remote service requests**. Remote service is disabled by default. Disabling remote service does not prevent the console from detecting and analyzing problems. The console still issues hardware messages to notify console operators of problems.

b. **Authorize automatic service call reporting**. If this feature is selected, the console automatically reports problems and requests service.

   If service is required to correct a problem, the hardware message includes instructions for calling a customer service center to report the problem and request service.

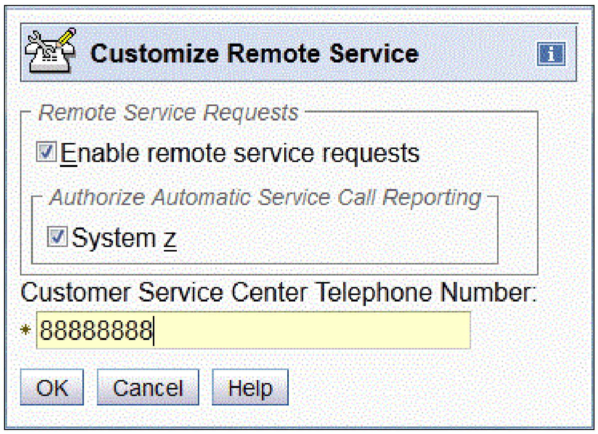

Figure 3-25 HMC and SE: Customize Remote Services

3. Select **Customize Outbound Connectivity** → **Configure** and click the **Internet** tab. Select **Enable the local console as a call-home server** and click **Test**.

   – To enable using an SSL proxy to access the Internet, select Use SSL proxy. If you select **Use SSL Proxy Connection to Internet**, provide an Internet Protocol address and port number to direct requests. The Internet Protocol address can be specified as either a host name or TCP/IP address (IPv4 or IPv6 format).

   – If your proxy requires authentication to forward requests, select **Use SSL Proxy Authenticate** and enter the user name and password. This name and password are used for all requests that are called home from this HMC. See Figure 3-26 on page 74.

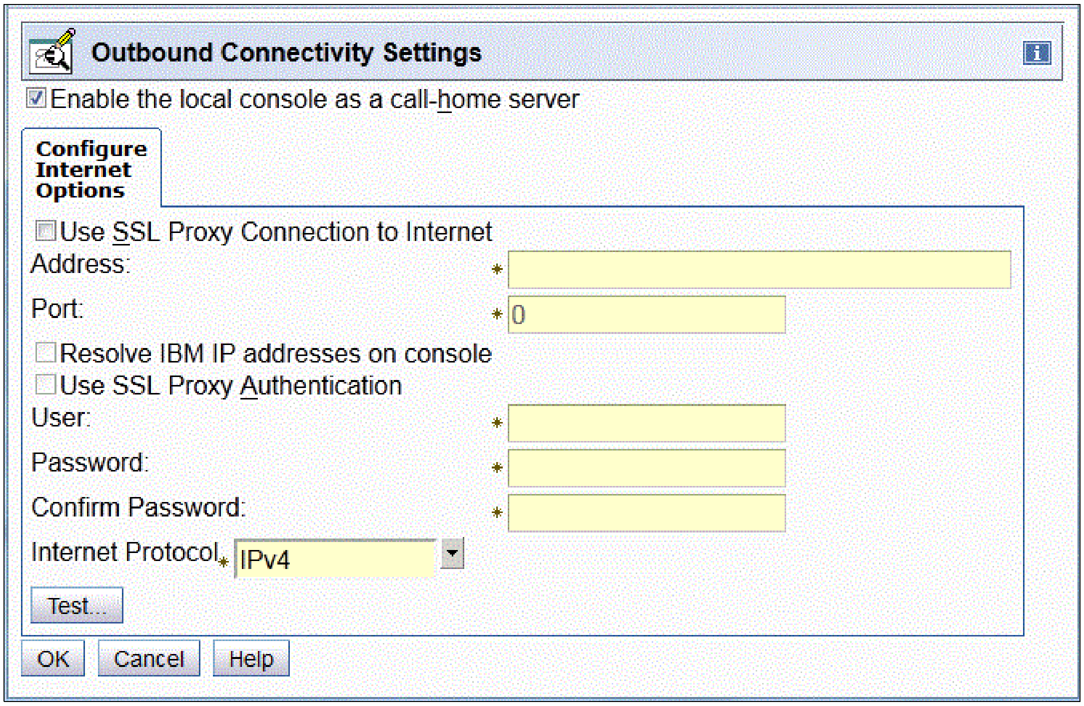

Figure 3-26 Outbound Connectivity Settings

4. The Test Internet window opens. Click **Start** and wait for the results.
5. Select **Customize Network Setting** from HMC management to open the Customize Network Settings window (Figure 3-27).

> **Requirement:** To access the Customize Network Settings window, you must be logged on with the *Access Administrator* role.

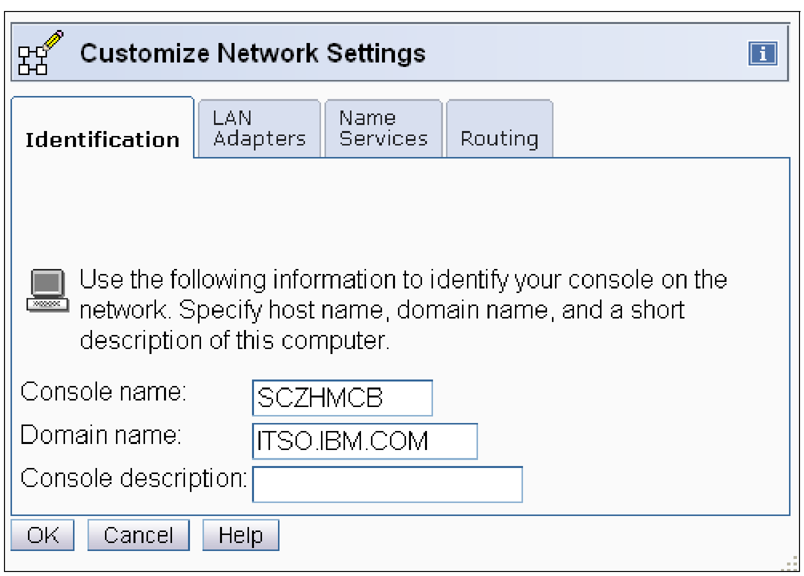

Figure 3-27 HMC and SE: Customize Network Setting

6. In the **Identification** tab, enter the console name, domain name, and a brief description.

7. Select the **LAN Adapters** tab. Select one of the adapters.
8. Click **Details** to open the LAN Adapter Details window.
9. Select the IPv4 or IPv6 setting, make the required entries, and click **OK**. Repeat this procedure for the second LAN Adapter if necessary. For more information, see "TCP/IP settings on the HMC and SE" on page 86.
10. Click the **Name Services** tab in the Customize Network Settings window and complete these steps:
    a. Check the **DNS enabled** check box, if required.
    b. Enter the DNS Server Search Order IP address, if required, and click **Add**.
    c. Enter the Domain Suffix Search Order, if required, and click **Add**.
11. Select the **Routing** tab in the Customize Network Setting window.

    This information is needed to connect to RETAIN through the Internet. Select **Enable 'routed'** only when there is a request. Click **New** to add the routing information.
12. Make the appropriate selections and entries on the Route Entry window:
    - Route Type
    - Destination, Gateway, and Subnet mask to match your network requirements
    - Adapter
13. Click **OK**. You are taken back to the Customize Network Settings window. Enter the Gateway information and click **OK**.

A message indicates that the network settings are successfully updated.

Return to step 3 on page 73 to check your changed network settings. Remember that you must leave the Access Administrator role to go back to the Customize Outbound Connectivity function.

3.2.8 Capacity on demand

To fulfill unpredictable client needs, market opportunities, and external pressure without interrupting existing processes, your IT infrastructure must support changing business objectives. You should have access to the resources you require only when you need them.

This is the basic principle that underlies the Capacity on Demand (CoD) offerings for the z13. The CoD offerings allow you to perform permanent upgrades and temporary upgrades to a z13 server.

A permanent upgrade can provide these services:

- Increase model capacity
- Increase the number of specialty engines:
  - Integrated Coupling Facilities (ICFs)
  - Integrated Facilities for Linux (IFLs)
  - System assist processors
- Increase memory
- Add channels
- Add cryptos
- Change specialty engines (recharacterization)

A temporary upgrade can provide these services on a temporary basis:

- Increase model capacity
- Increase the number of specialty engines:
  - Integrated Coupling Facilities (ICFs)
  - Integrated Facilities for Linux (IFLs)
  - System assist processors

New in SE version 2.13.0: Some panels that had preselected defaults are removed. You must manually specify each selection.

Use one of the following methods to retrieve and apply (install) a permanent upgrade record as shown in Figure 3-28 on page 77:

- To retrieve and immediately apply the permanent upgrade record, complete these steps:
  a. From the Perform Model Conversion window, click **Permanent upgrades**, **Retrieve and apply**, then **Processor/memory upgrade data from IBM Service Support System**.
  b. On the Customer Initiated Upgrade Order Activation Number window, enter the order activation number and click **OK**.
- To retrieve the permanent upgrade record and apply it later, complete these steps:
  a. From the Perform Model Conversion window, click **Permanent upgrades** then **Retrieve processor/memory upgrade data but do not apply**.
  b. In the Customer Initiated Upgrade Order Activation Number window, enter the order activation number and click **OK**.

     The upgrade record is retrieved and stored on the Support Element hard disk drive to be installed later.

     If you decide you do not want to apply the upgrade record, contact IBM support who will send someone to remove the record.
  c. When you are ready to apply the upgrade record you previously retrieved *and* no temporary upgrade records are activated, the **Perform precheck on staged upgrade** tasks displays in the Perform Model Conversion window. However, it is disabled. This means that a permanent upgrade record is staged, but no temporary upgrade records are activated. Therefore, the precheck function is not necessary.
  d. When you are ready to apply the upgrade record you previously retrieved *and* temporary upgrade records are activated **OR** if you must verify that a previous conflict is corrected, the **Perform precheck on staged upgrade** task displays in the **Perform Model Conversion** window. This process allows you to check for any conflicts before you apply the permanent upgrade record.

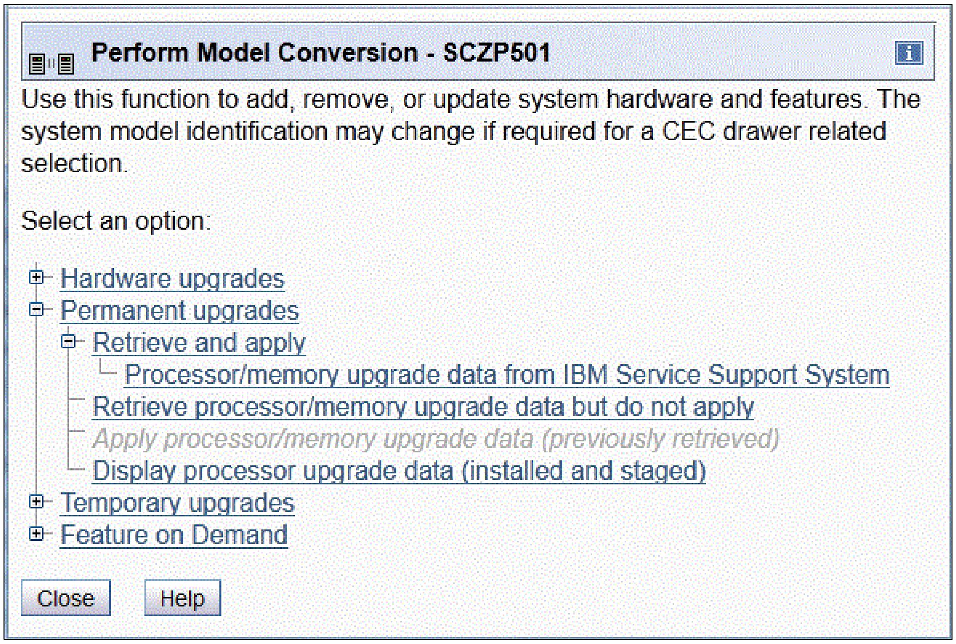

Figure 3-28 Perform Model Conversion, permanent upgrades

Temporary upgrade

You can activate a temporary upgrade record by using any of the following methods:

- Manually activating an upgrade record by using the Support Element
- Setting up scheduled operations to activate upgrade records
- Using APIs
- z/OS Capacity Provisioning (*z/OS MVS Capacity Provisioning User's Guide* , SC33-8299*)*

For details, see the *zEnterprise System Capacity on Demand User's Guide,* SC28-2605.

3.2.9 Server Time Protocol (STP) enhancements

Enhancements and changes in z13 regarding STP are described in this section.

ETR removal

Note: The ETR function is removed from the z13. The ETR function now relies solely on STP for time STP enhancements. A mixed CTN is no longer supported for z13 and later.

Preventing inadvertent changes

A requirement was fulfilled to make it harder for customers to inadvertently make changes to local offsets and time zone incorrectly (see Figure 3-29 on page 78):

- Added confirmation messages when setting STP time zone via Adjust Time Zone panel on the Current Time Server (CTS)
- Listing scheduled switch times for leap seconds and time-zone/daylight saving time on the Timing Network tab

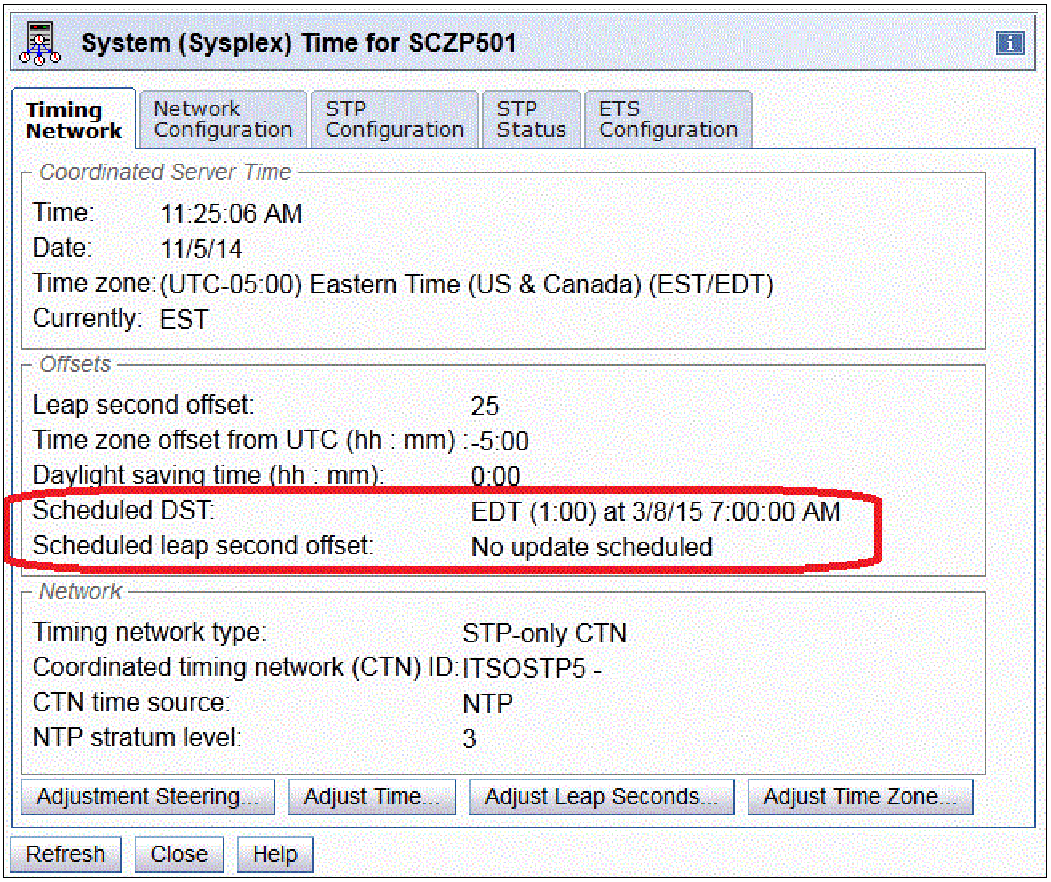

Figure 3-29 Sysplex Time Panel

View-only panels

A view-only option was added to give more control as to who can make changes to STP functionality. By changing an existing role or adding a new user role on an HMC, an administrator can limit specific users to view STP settings only. To implement the view-only role, perform the following steps:

1. Create new user task role by copying an existing sysprog role as a template. Select **System (SYSPLEX) Time** on the right-side panel, under Current Tasks, and click **Remove**. Then, select **System (SYSPLEX) Time** on the left-side panel under Available Tasks and click **Add**. See Figure 3-30 on page 79.

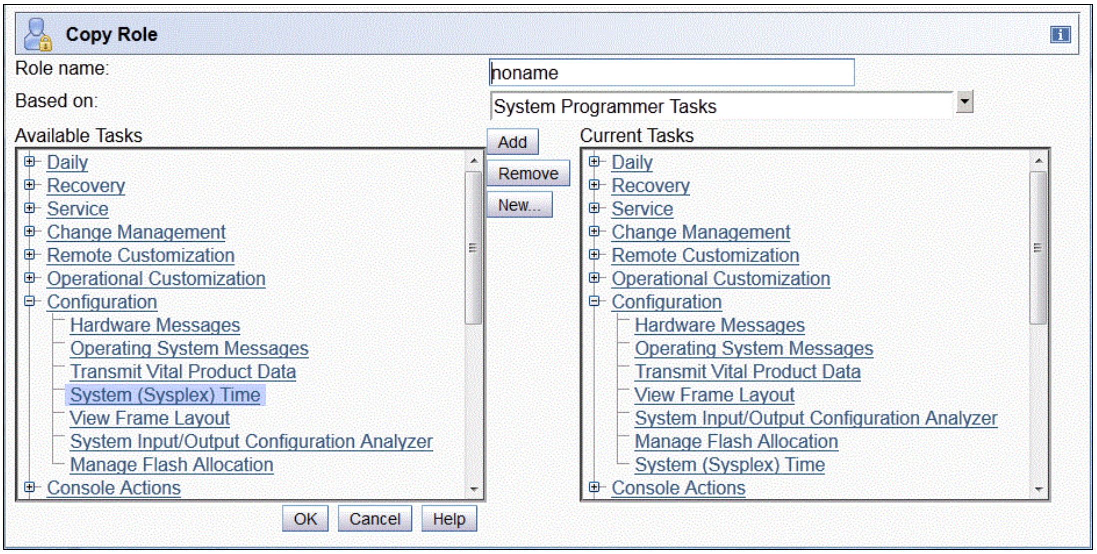

Figure 3-30 HMC User Role Modification panel

You are presented with the option to allow a view-only option for the Sysplex Time function (Figure 3-31). If you select **No**, the task will have full STP control.

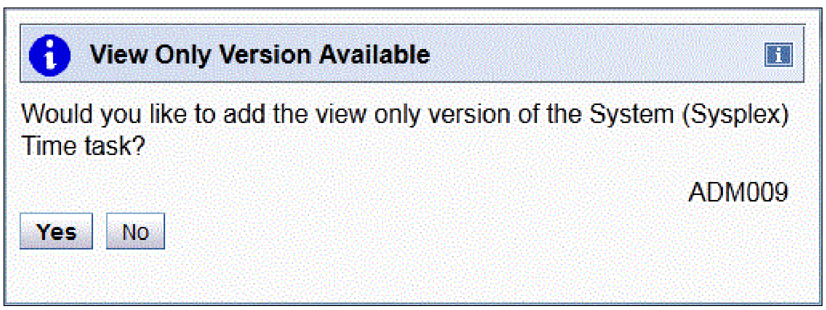

Figure 3-31 Sysplex Time view only option

2. Create or modify a user with that new task role to give attribute to the user.

Moved: Enable for time synchronization check box

The Enable for time synchronization check box was removed from the Add or Change Object Definition task. The Add or Change Object panel before the check box was removed is shown in Figure 3-32 on page 80; the panel after the check box was removed is shown in Figure 3-33 on page 80.

> **Note:** To view or alter the information on the Add or Change Object panel, you must be logged on as *Access Administrator*.

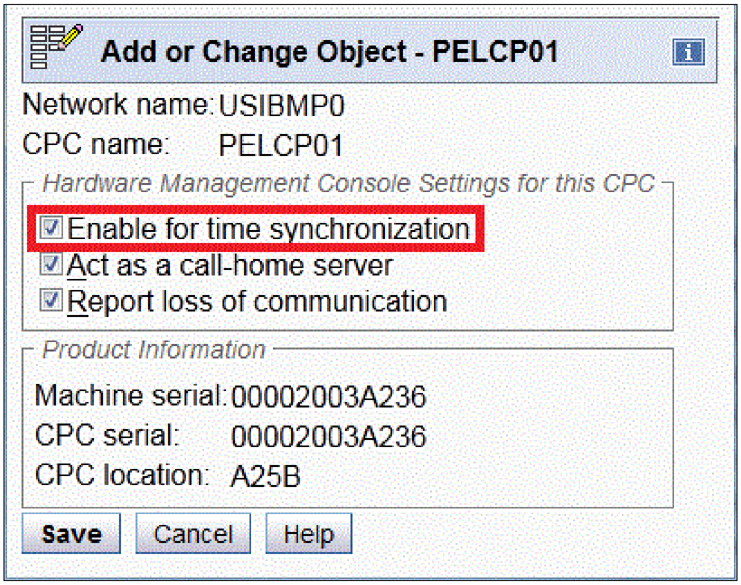

Figure 3-32 Pre z13 Add or Change Object

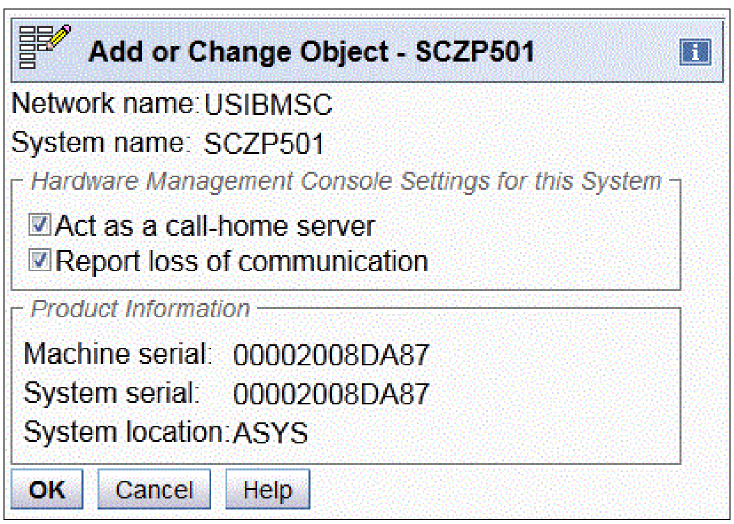

Figure 3-33 z13 Add or Change Object

The Selected CPCs option as a time source selection was moved to the Customize Console Date/Time task (see Figure 3-34 on page 81).

Note: If NTP is used by the HMC, this enhancement is not relevant.

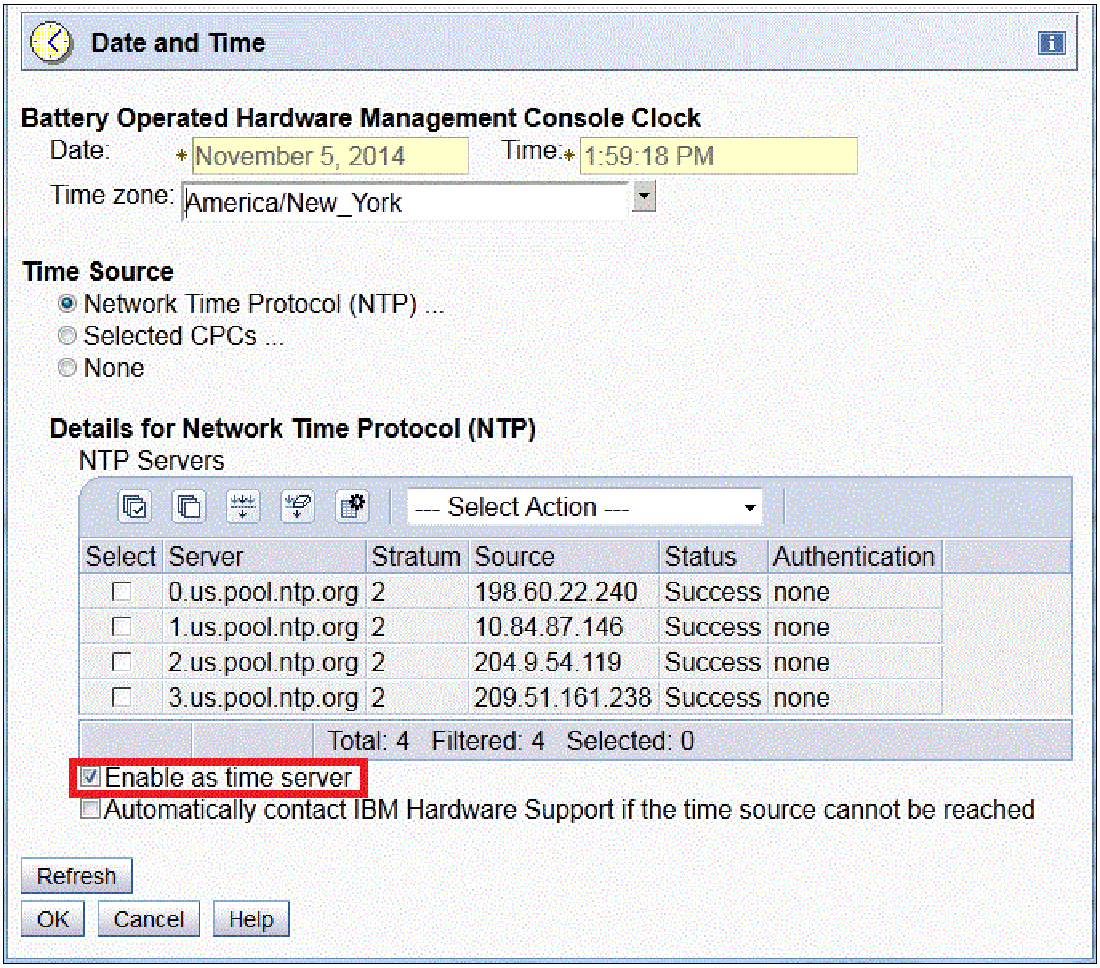

Figure 3-34 Customize Console Date and Time

Redesigned: Customized Console Date and Time panel

The panel was redesigned; the following items were added:

- Clear selection of HMC time source
- Manual setting of date and time only when Time Source is set to **None**.

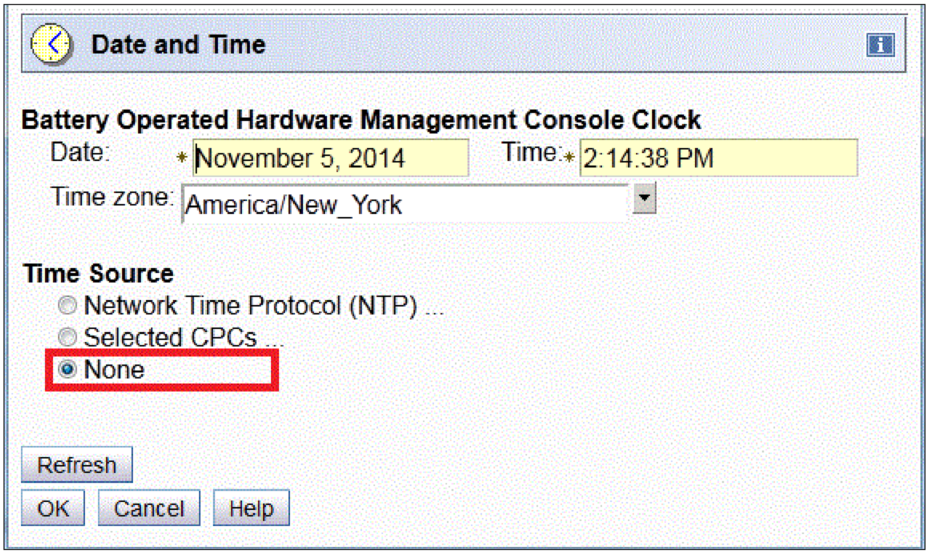

Figure 3-35 Customized Console Date and Time Panel

Added: Configuration support for time retrieval from selected CPCs

One location for configuration and checking all CPCs used as a time source. Additional checks added to promote homogenous time sources when using multiple CPCs as time sources (same STP CTN). A warning is issued if the same Coordinated Timing Network (CTN) sources are not selected.

CPC selected for Add or Remove Object Definition task

A panel displayed to allow adding or removing of selected CPCs as a time source if HMC is not configured for NTP (see Figure 3-36).

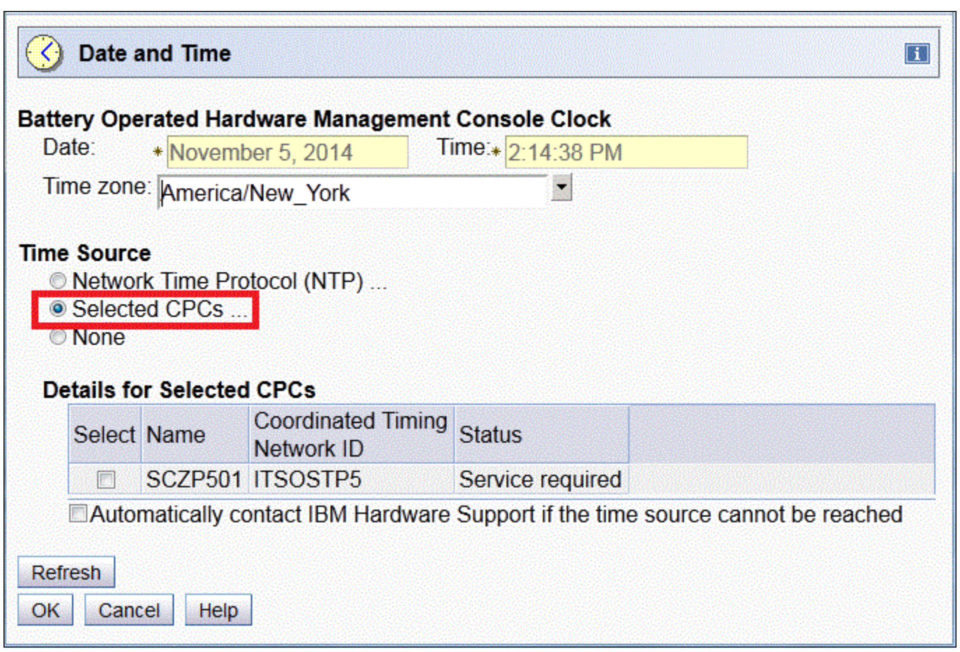

Figure 3-36 STP CPC Time Source selection panel

A confirmation panel opens when a CPC time source is added or removed (Figure 3-37).

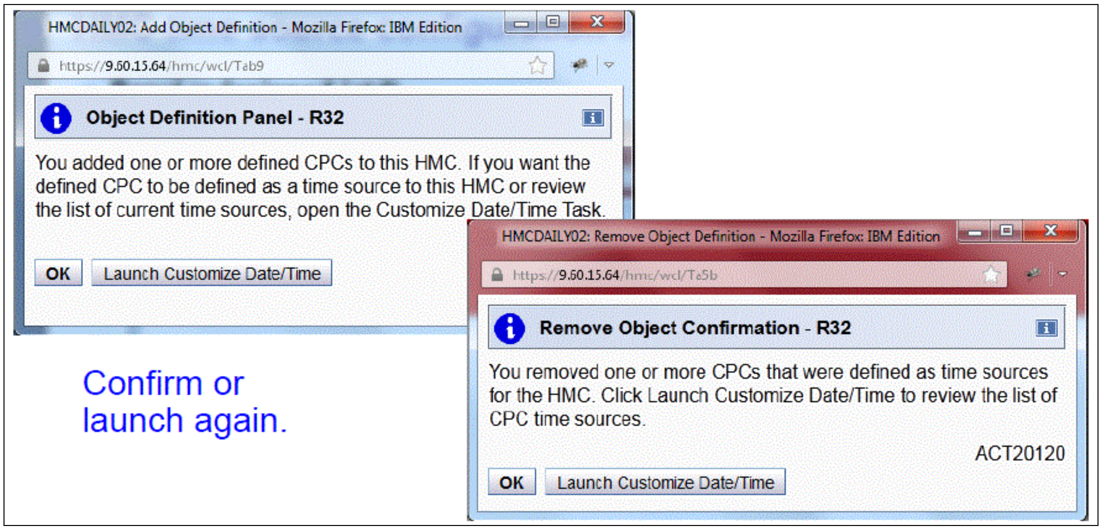

Figure 3-37 Confirming Object addition or removal

NTP thresholds

NTP threshold capability function was added with zEC12. You can use this function to specify threshold settings to suppress generation of hardware and operating system messages related to changes in the NTP server stratum level or source ID. Operating system messages are only generated if the operating system supports posting of messages to notify clients of STP-related hardware messages.

To configure the NTP thresholds, complete the following steps:

1. In System (Sysplex) Time task, click the **ETS Configuration** tab (Figure 3-38) and click **NTP Thresholds**.

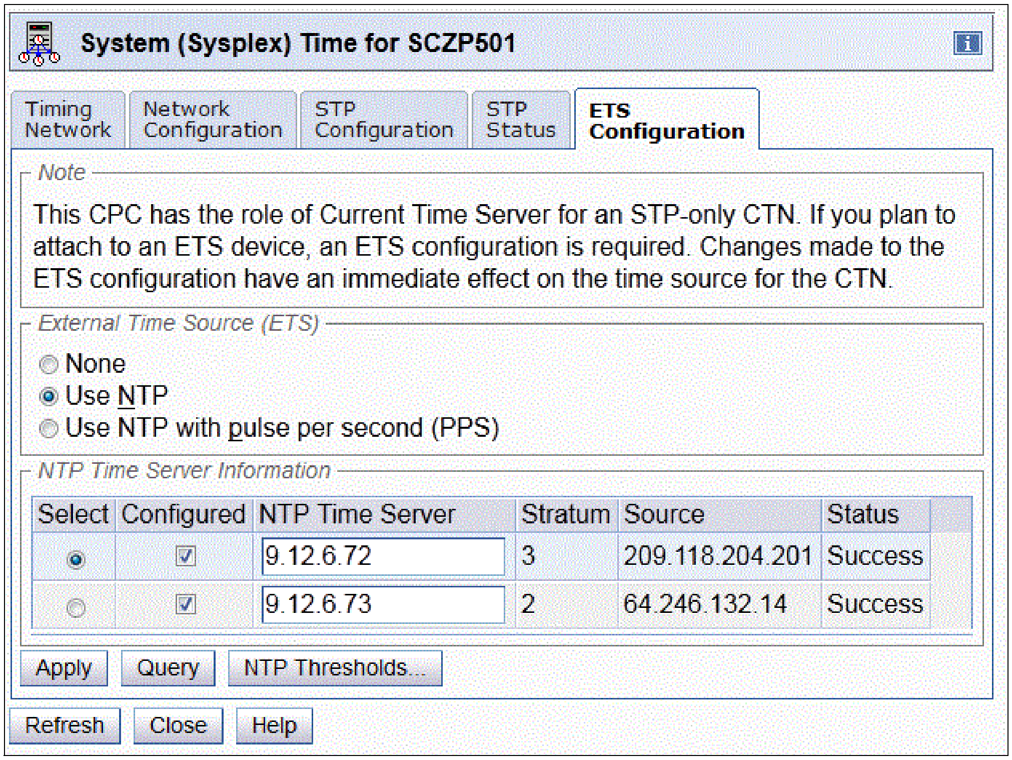

Figure 3-38 ETS configuration

2. In the NTP Thresholds window (Figure 3-39), you can change the following variables:
   - **Stratum level threshold**: Indicates the NTP server stratum level that must be reached before the system generates a hardware and operating system message. The threshold can be set as low as 2 and as high as 15.
   - **Source ID time threshold**: Indicates the amount of time that must pass before a change in the source ID generates a hardware and operating system message. The messages are issued if the source ID does not return to the original value within the specific time period. The threshold can be set as low as 0 and as high as 24 hours, in increments of one half hour.

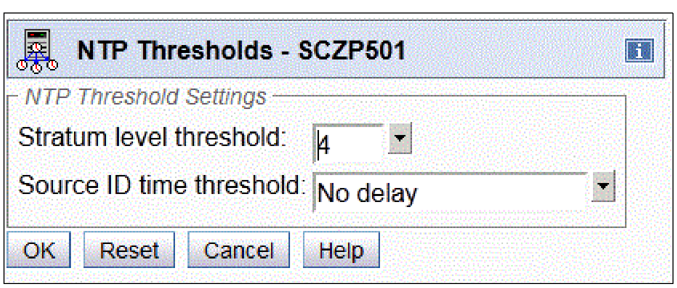

Figure 3-39 NTP Thresholds

The HMC can still be used to connect a network time protocol (NTP) server to maintain time accuracy for the connected CPCs. Details about the setup of the STP are in Chapter 9, “Server Time Protocol (STP)” on page 511.

3.2.10 Connectivity for the HMC and the SE without ensemble enabled

Network connectivity options are available for the HMC and SE when the z13 is *not* part of an ensemble.

Hardware connectivity

The HMC has two Ethernet adapters, which can be configured to connect to two different Ethernet LANs. It is no longer possible to order Ethernet switches that are required by the HMCs to connect to the z13. You must provide new or existing Ethernet switches.

The Ethernet switch/hub generally has these characteristics:

- 16 auto-negotiation ports
- 10/100/1000 Mbps data rate
- Full or half duplex operation
- Auto-MDIX on all ports
- Port Status LEDs

The SEs on the z13 are connected to the System Control Hub (SCH). The HMC to SCH communication is possible only through an Ethernet switch. Other z Systems and the HMCs can also be connected to the switch. To provide redundancy for the HMCs, use two switches.

> **Note:** On zEnterprise CPCs (zEC12, z196), the SEs were connected to a Bulk Power Head (BPH).

Only the switch (and not the HMC directly) can be connected to the SCH ports J01 and J02 for the client network 1 and 2. The SE is always connected to SCH port J05.

Figure 3-40 shows the connectivity between the HMC and the SE.

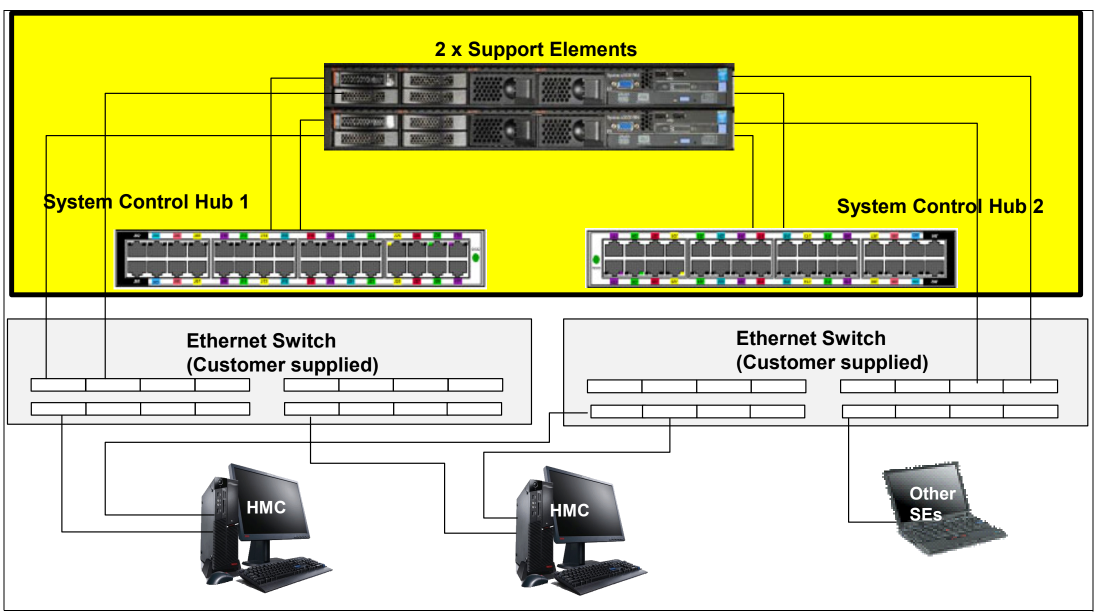

Figure 3-40 HMC to SE connectivity

Various methods are available for setting up the network. You must plan and conceive the HMC and SE connectivity. Select the method based on your environment's connectivity and security requirements.

Plan the HMC and SE network connectivity carefully to allow for current and future use. Many z Systems capabilities benefit from the various network connectivity options available. For example, the following functions are available to the HMC that depend on the HMC connectivity:

- LDAP support that can be used for HMC user authentication
- NTP client/server support
- RSF through broadband
- The HMC remote web browser
- Enablement of the SNMP and CIM APIs to support automation or management applications such as IBM System Director AEM

These examples are shown in Figure 3-41.

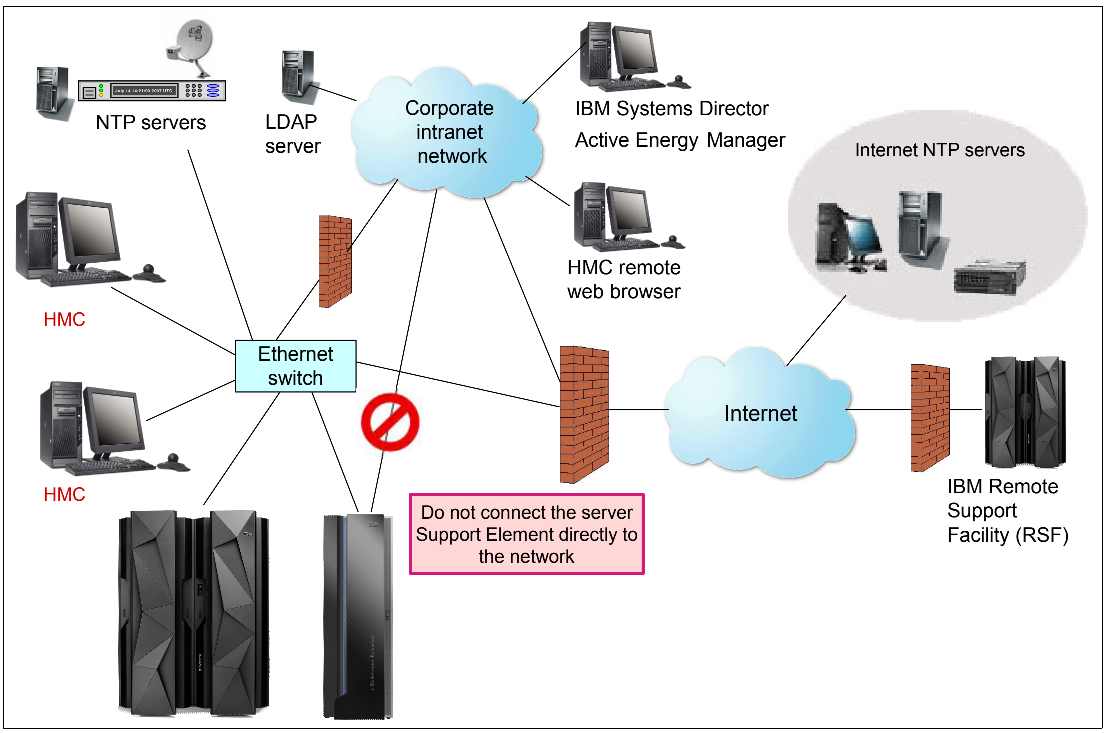

Figure 3-41 The HMC and SE: HMC and SE Connection without ensemble

For more information, see the following publications:

- *Hardware Management Console Operations Guide (V2.12.0),* SC28-6919
- *Installation Manual - Physical Planning 2827,* GC28-6914

TCP/IP settings on the HMC and SE

The HMC and SE can communicate by using IPv4, IPv6, or both. Assigning a static IP address to an SE is unnecessary if the SE has to communicate with only the HMCs on the same subnet.

To configure the HMC and SE IP settings, use the **Customize Network Settings** task to configure LAN adapter, name services, and routing (Figure 3-42 on page 87).

Note: By default, these tasks are available only to the ACSADMIN and SERVICE users.

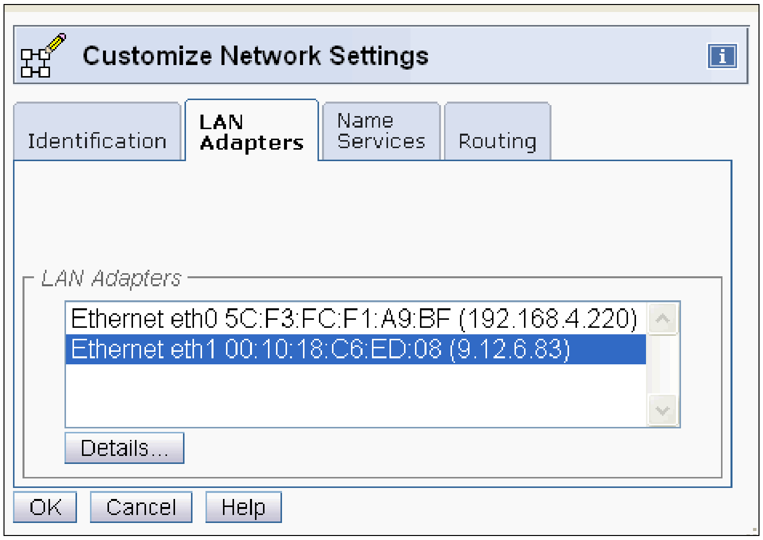

Figure 3-42 Customize Network Settings on the HMC

To configure the IP addresses, open the LAN Adapters tab, and click **Details** for each LAN adapter.

The HMC can have the following IP configurations:

- Statically assigned IPv4 or statically assigned IPv6 addresses.
- DHCP assigned IPv4 or DHCP assigned IPv6 addresses.
- Autoconfigured IPv6:
  - Link-local is assigned to every network interface.
  - Router-advertised, which is broadcast from the router, can be combined with a media access control (MAC) address to create a unique address.
  - Privacy extensions can be enabled for these addresses as a way to avoid using MAC address as part of address to ensure uniqueness.

An SE can have the following IP addresses:

- Statically assigned IPv4 or statically assigned IPv6.
- Autoconfigured IPv6 as link-local or router-advertised.

SE IP configuration cannot be dynamically assigned through DHCP to ensure repeatable address assignments. Privacy extensions are not used.

The HMC uses IPv4 and IPv6 multicasting to automatically discover SEs. The **HMC Network Diagnostic Information** task (Figure 3-43) can identify the IP addresses (IPv4 and IPv6) being used by the HMC to communicate to the CPC SEs.

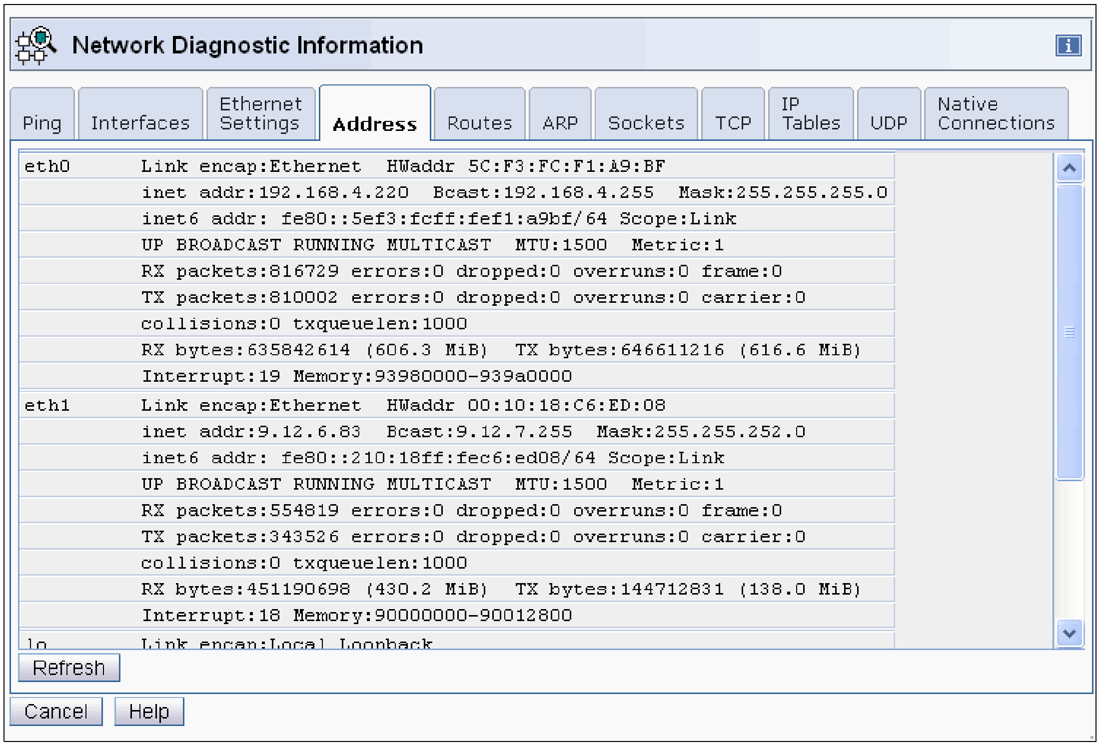

Figure 3-43 Network Diagnostic Information

IPv6 address on the HMC and SE

The HMC and SE can use IPv6 link-local addresses to communicate with each other.

IPv6 link-local addresses have the following characteristics:

- Every IPv6 network interface is assigned a link-local IP address.
- A link-local address is for use on a single link (subnet) and is never routed.
- Two IPv6-capable hosts on a subnet can communicate by using link-local addresses, without having any other IP addresses (v4) assigned.

IPv6 addresses are easily identified. A fully qualified IPV6 address has 16 bytes, written as eight 16-bit hexadecimal blocks separated by colons (:), as shown in this example:

```
2001:0db8:0000:0000:0202:b3ff:fe1e:8329
```

Because many IPv6 addresses are not fully qualified, shorthand notation can be used. Leading zeros can be omitted, and a series of consecutive zeros can be replaced with a double colon (`::`). The address in the previous example can also be written as follows:

```
2001:db8::202:b3ff:fe1e:8329
```

For remote operations that use a web browser, if an IPv6 address is assigned to the HMC, navigate to it by specifying that address. The address must be enclosed in square brackets in the browser's address field:

```
https://[fdab:1b89:fc07:1:201:6cff:fe72:ba7c]
```

Using link-local addresses must be supported by your browser.

Network security

Network security is another important concern. The HMC must be attached to a network so that it can manage the z Systems resources associated with it. In some cases, for the HMCs located close to the z Systems servers it is managing, this network is a *private* network that is fully contained on a single raised floor. However, when a client has multiple data centers or attaches the HMC to a corporate intranet to allow for remote access, network security is of utmost importance.

Because the HMC can be a multi-homed system (that is, it has multiple network interfaces), it can be connected to a private network that contains the z Systems resources and the corporate intranet at the same time. This is a common client configuration because it provides a level of physical separation for the z Systems resources, while simultaneously allowing for the use of advanced HMC capabilities. These capabilities include remote access and Internet connectivity for remote support.

The HMC Licensed Internal Code includes a full-function firewall that controls network access to the HMC. By default, the HMC allows virtually no inbound network traffic. As different features of the HMC are enabled, such as remote access and SNMP-based automation, more inbound network traffic is allowed. Table 3-5 lists the various TCP/IP ports that the HMC uses for *inbound* network traffic.

Table 3-5 TCP/IP port uses for inbound network

TCP/IP source port	Use
ICMP Type 8	Ping to and from the HMC and the z Systems resources the HMC manages.
tcp 58787	Automatic discovery of z Systems servers.
tcp 4455	Automatic discovery of Director/Timer console.
udp 9900	HMC-to-HMC automatic discovery.
tcp 55555	SSL-encrypted communications to and from z Systems servers. The internal firewall allows only inbound traffic from z Systems servers that are defined to the HMC.
tcp 9920	HMC-to-HMC communications.
tcp 443	Remote user access to the HMC. Inbound traffic for this port is allowed only if remote access is enabled for the HMC.
tcp 9950-9954	Proxy **Single Object Operations** sessions for a z Systems server.
tcp 9960	Remote user applet-based tasks. Inbound traffic for this port is allowed only if remote access is enabled for the HMC.
tcp 21	Inbound FTP requests. This port is enabled *only* when an IBM Electronic Service Agent™ or the **Enable FTP Access to Hardware Management Console Mass Storage Media** task is in use. FTP is an unencrypted protocol so, for maximum security, do not perform these tasks on the HMC.
udp/tcp 161	SNMP automation. Inbound traffic for these ports is allowed only when SNMP automation is enabled.
tcp 5988 tcp 5989	CIM automation. Inbound traffic for these ports is allowed only when CIM automation is enabled.
tcp 6794	Web services SSL encrypted automation traffic. Inbound traffic for this port is allowed only when Web Services automation is enabled.

Outbound security

The HMC also initiates requests to the z Systems resources that it is managing and to the other HMCs. Table 3-6 lists the types of *outbound* network traffic that the HMC initiates.

Table 3-6 . TCP/IP port uses for outbound network

TCP/IP destination port	Use
ICMP Type 8	Ping to and from the HMC and the z Systems resources the HMC manages.
udp 9900	HMC-to-HMC automatic discovery.
tcp/udp 58787	Automatic discovery and establishment of communications with z Systems servers.
tcp 55555	SSL-encrypted communications to and from z Systems servers. The internal firewall allows only inbound traffic from the z Systems servers that are defined to the HMC.
tcp 9920	HMC-to-HMC communications.
tcp 443	Single Object Operations to a z Systems server console.
tcp 9960	Proxy remote user applet-based tasks during a Single Object Operations session for a z Systems server console.
tcp 25345	Single Object Operations session to an older z Systems server console.
tcp 4455	Communicating with Director/Timer consoles the HMC manages.
udp 161	Communicating with an IBM Fiber Saver that the HMC manages.
tcp 25	Send email events to an SMTP server for delivery by using the Monitor System Events task, when the HMC is configured. This can be a port other than 25, but 25 is the default SMTP port that most SMTP servers use.

Browser-based security

Because access to the HMC is provided over an HTML-based user interface, current browser-based encryption mechanisms are employed for enhanced security. All remote browser access to the HMC must use SSL encryption. Only the local user interface can use non-encrypted access because that is inherently a secure environment. With SSL encryption required for all remote access to the HMC, a certificate is required to provide the keys that are used for this encryption. By default, the HMC provides a self-signed certificate that allows this encryption to occur.

For more information about HMC security, see the “IBM System z HMC Security” topic, available at the IBM Resource Link[2] website:

https://www-304.ibm.com/servers/resourcelink/lib03011.nsf/pages/zHmcSecurity/$file/zHMCSecurity.pdf

For more information about access and security on the HMC, see *Hardware Management Console Operations Guide (V2.12.0),* SC28-6919.

[2] To access the IBM Resource Link, registration is required.

3.2.11 Connectivity for the HMC and the SE with ensemble enabled

An ensemble is a platform systems management domain that consists of up to eight z13 or zEnterprise nodes. A node can be a z13, zEnterprise CPC or a zBX. The ensemble provides an integrated way to manage virtual server resources and the workloads that can be deployed on those resources. The z13 is a workload optimized technology system that delivers a multi-platform, integrated hardware system. This system spans z Systems, IBM System p®, and IBM System x blade server technologies.

Management of the ensemble is provided by the IBM zEnterprise Unified Resource Manager.

When you set up an ensemble, start with a pair of HMCs and assign an ensemble identity. For more information, see Chapter 12, “Preparing IBM z13 for zBX Model 004” on page 595.

Primary and alternate HMC

If a CPC is part of an ensemble, the CPC Details task on the SE and the HMC reflect this situation in the ensemble name. Two HMCs are required within the ensemble for high availability ensemble management. The HMC that creates an ensemble (that is. the Create Ensemble wizard) becomes the primary HMC. An alternate HMC is then selected and paired with the primary HMC.

Unified Resource Manager actions for the ensemble are conducted from a single, primary HMC. All other HMCs connected to the ensemble are able to run system management tasks (but not ensemble management tasks) for any CPC within the ensemble. The primary HMC can also be used to run system management tasks on CPCs that are not part of the ensemble.

The ensemble-specific managed objects include these items:

- Ensemble
- Members
- Blades
- BladeCenters
- Hypervisors
- Storage resources
- Virtual servers
- Workloads

Other HMC

When another HMC accesses an ensemble node’s CPC, the HMC can do the same tasks as if the CPC were not a part of an ensemble. Some tasks have been extended so you can configure certain ensemble-specific properties such as setting the virtual network associated with OSAs for an LPAR. Showing ensemble-related data in certain tasks is allowed. Generally, if the data affects the operation of the ensemble, the data is read-only on another HMC. The following tasks can show ensemble-related data on another HMC:

- Scheduled operations: Displays ensemble introduced scheduled operations, but you can view only these scheduled operations.
- User role: Shows ensemble tasks, and you can modify and delete those roles.
- Event monitoring: Displays ensemble-related events, but you cannot change or delete the event.

HMC considerations when used to manage an ensemble

Consider the following information for using Unified Resource Manager to manage an ensemble:

- All HMCs at the supported code level are eligible to create an ensemble. Only HMCs with Feature Code 0094 and Feature Code 0092 are capable of being the primary or alternate HMCs for a z13 server.

> **Note:** Older HMCs, FC 0091 and FC 0092, can be upgraded to control a z13.

- A single HMC pair manages the ensemble: primary HMC and alternate HMC. The primary and the alternate HMC must be the same machine type/feature code.
- Only one primary HMC manages an ensemble which can consist of a maximum of eight CPCs.
- The HMC that ran the Create Ensemble wizard becomes the primary HMC. An alternate HMC is elected and paired with the primary.
- "Primary Hardware Management Console (Version 2.13.0 or later)" and "Alternate Hardware Management Console (Version 2.13.0 or later)" are displayed on the HMC banner. When the ensemble is deleted, the titles resort to the default.
- A primary HMC is the only HMC that can run ensemble-related management tasks such as create virtual server, manage virtual networks, and create workload.
- A zEnterprise ensemble can have a maximum of eight nodes and is managed by one primary HMC and its alternate. A node can be a z13 CPC, zEnterprise CPC or an IBM z BladeCenter Extension (zBX Model 004).
- Any HMC can manage up to 100 CPCs. The primary HMC can run all non-ensemble HMC functions on CPCs that are not members of the ensemble.
- The primary and alternate HMCs *must be on the same LAN segment*.
- The alternate role of the HMC is to mirror ensemble configuration and policy information from the primary HMC.
- When failover happens, the alternate HMC becomes the primary HMC. This behavior is the same as the current primary and alternate Support Elements.

Requirements of z13 for participation in an ensemble

The requirements are as follows:

- Ensemble and quality of service (QoS) features:
  - Ensemble Feature, FC 0025, is always required for ensemble participation.
  - Choose a QoS selection: either FC 0019-Manage level or FC 0020-Automate level.

> **Note:** All nodes in the same ensemble *must* have the same QoS feature level.

- Intranode management network (INMN) connectivity is always required. Two OSA-Express 1000BASE-T features are needed to support the two required OSM CHPIDs. (OSA-Express5S FC 0417 or OSA-Express4S FC 0408: Carry Forward only). Two TYPE=OSM CHPIDs must be defined for each of the above and are cabled to the z13 internal System Control Hub.
- If one or more optional zBX nodes are to be communicated with, an intraensemble data network (IEDN) connectivity is required. One or more pairs of OSA-Express 10GbE features are required for the data traffic. (OSA-Express5S 10 GbE LR FC 0415 or OSA-Express4S 10 GbE LR FC 0406, OSA-Express5S 10 GbE SR FC 0416 or

OSA-Express4S 10 GbE SR FC 0407). Two TYPE=OSX CHPIDs must be defined for each of the above pairs and are cabled to a common customer network.

Note: Be sure that SFP optics features ordered for zBX match (SR versus LR) and are cabled to the matching optics in the IEDN TOR switches in zBX.

- Hardware for both ensemble primary and alternate HMCs must be identical. Two HMC FCs 0091 or 0092 (deskside) or 0094 (rack mount) and both HMCs must be at Driver Level 22.

Note: At Driver Level 22, the Ensemble HMCs also support nodes including zEC12, zBC12, z196 and z114 with or without managed zBX Model 3 or Model 2.

HMC availability

The HMC is attached to the same subnet as the server's SE. This LAN is referred to as the *Customer Managed Management Network*. The HMC communicates with each CPC and zBX throughout the network.

If the z13 node is defined as a member of an ensemble, the primary HMC is the authoritative controlling component for Unified Resource Manager configuration and policies. The scope of control spans all of the managed CPCs/SEs in the ensemble. The managing HMC has an active role in ongoing system monitoring and adjustment. This role requires that the HMC is configured in an primary/alternate configuration and cannot be disconnected from the managed ensemble members.

Failover: The primary HMC and its alternate must be connected to the same LAN segment. This configuration allows the alternate HMC to take over the IP address of the primary HMC during failover processing.

Considerations for multiple HMCs

Clients often deployed multiple HMC instances to manage an overlapping collection of systems. Until the emergence of ensembles, all of the HMCs were peer consoles to the managed systems. In this configuration, all management actions are possible to any of the reachable systems while logged in to a session on any of the HMCs (subject to access control). With the Unified Resource Manager, this paradigm has changed. One ensemble is managed by one primary and alternate HMC pair. Multiple ensembles require an equal number of multiple primary and alternate HMC pairs to manage them. If a z13 or zEnterprise System node is added to an ensemble, management actions that target that system can be done only from the managing (primary) HMC for that ensemble.

The HMC browser session to a primary HMC

A remote HMC browser session to the primary HMC that is managing an ensemble allows a user who is logged on to another HMC or a workstation to perform ensemble-related actions.

3.3 Channel considerations

With the introduction of PCIe-O SR coupling links for the z13 server, new definitions and rules were introduced. The planning necessary to implement them is covered in 3.3.6, “Integrated Coupling Adapter (PCIe-O SR) Adapter” on page 103.

Certain types of channels are no longer supported on the z13. Some feature codes can be carried over to the z13 but are no longer orderable features. If you have critical devices on a channel type that is no longer supported, make alternative arrangements for connectivity before you upgrade to the z13 server.

This section introduces the connectivity options available on a z13 server. Table 3-7 lists the channel types and definitions that are supported on a z13. Details about system-relevant rules for z13 are described in the sections that follow.

Table 3-7 Supported channel types and definitions

Channel type	CHPID definition	MIF[a]-shared	Spanned
FICON Express16S	FC and FCP	Yes	Yes
FICON Express8S[b]	FC and FCP	Yes	Yes
FICON Express8[c]	FC and FCP	Yes	Yes
OSA Express5S[b]	OSD, OSE, OSC, OSN, OSM, and OSX	Yes	Yes
OSA Express4S[c]	OSD, OSE, OSC, OSN, OSM, and OSX	Yes	Yes
PSIFB	CIB	Yes	Yes
ICA (PCIe-O SR)	CS5	Yes	Yes
IC (Internal CHPID)	ICP	Yes	Yes
HiperSockets (Internal CHPID)	IQD	Yes	Yes

a. multiple image facility (MIF)
b. This feature is available for new build or carry-forward.
c. This feature is carry-forward only.

3.3.1 Parallel channels

The z13 server does not support the direct attachment of parallel channels. Use a converter that converts ESCON connections to a parallel protocol enables devices with a parallel attachment.

Installing an FICON to ESCON converter allows you to attach an ESCON to Parallel channel converter which enables devices with a parallel attachment to be used (Figure 3-44 on page 95).

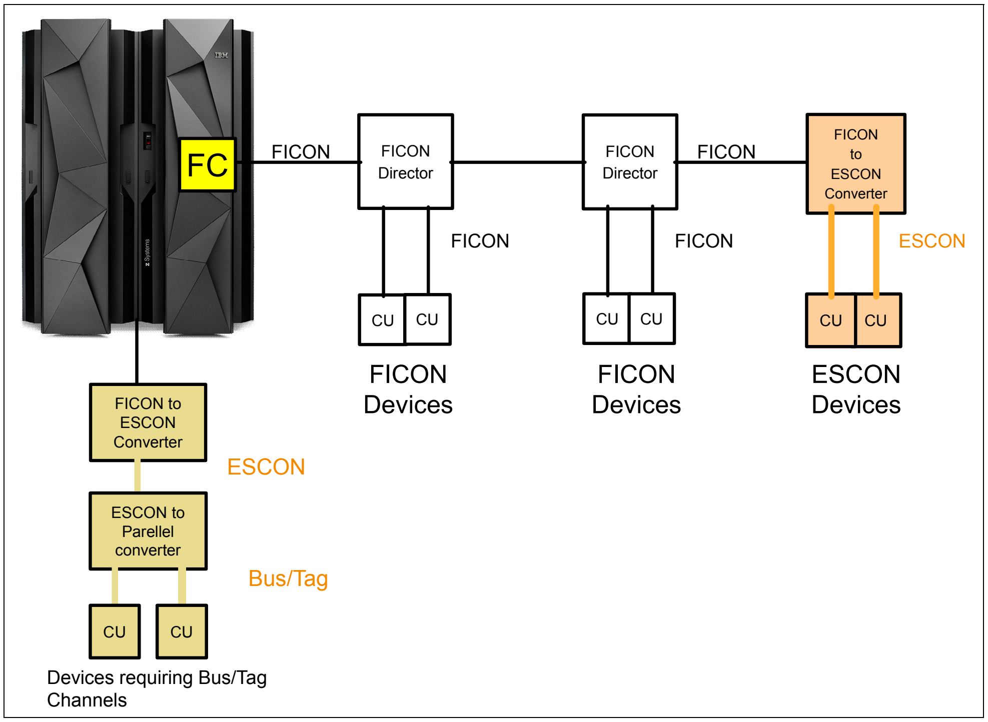

Figure 3-44 FICON to ESCON and ESCON to Parallel converter

3.3.2 ESCON channels and FICON bridge channels

z13 does not support ESCON Channels (CHPID CNC, CTC, CVC, CBY) and FICON bridge channels (CHPID FCV). If you have devices that need ESCON channels, use a FICON to ESCON converter. See Figure 3-44.

The PRIZM Protocol Converter Appliance from Optica Technologies Inc.® provides a FICON to ESCON conversion function that was qualified for z Systems. Qualification letters are at the I/O Connectivity website (under the Products tab in the FICON/FCP connectivity section):

http://www-03.ibm.com/systems/z/hardware/connectivity/products/index.html

For information about the PRIZM Protocol Converter Appliance, see Optica Technologies:

http://www.opticatech.com/

Attention: IBM cannot confirm the accuracy of compatibility, performance, or any other claims by vendors for products that have not been qualified for z Systems. Address questions about these capabilities and device support to the suppliers of those products.

IBM Facilities Cabling Services - ESCON to FICON migration services can help take advantage of high-speed FICON to support an upgraded IBM z Systems environment. Also, you can keep using ESCON-attached devices to reduce migration costs. For more information, see this web page:

http://www-935.ibm.com/services/us/index.wss/offering/its/c337386u66547p02

3.3.3 FICON channels

FICON Express16S and Express8S provides connectivity to storage devices by using Fibre Connection (FICON) or Fibre Channel Protocol (FCP). FICON Express16S supports auto negotiation of link data rate and supports 4 Gbps, 8 Gbps, and 16 Gbps link data rates. FICON Express8S supports auto negotiation of link data rate, and supports 2-Gbps, 4-Gbps and 8-Gbps link data rates. FICON Express16S and FICON Express8S support System z High Performance FICON (zHPF). zHPF is an extension to the FICON architecture that provides performance improvement for single-track and multi-track operations.

> **Note:** FICON Express16S features do not support a data rate of 2 Gbps. FICON Express8S features support 2 Gbps connectivity.

Previous servers had a 64 KB data transfer limit (that is, 16 x 4 KB) for a zHPF multi-track operation. From z196, the limit on the amount of data that can be transferred in a single operation was increased to 256 tracks. This new limit allows the channel to use the bandwidth of FICON Express16S and FICON Express8S channels when configured as CHPID type FC. The attached CU must support zHPF to use this enhancement.

A z13 server allows FICON channels to be defined in two ways:

- FICON native mode (FC)

  This type of channel allows for connectivity to control units that have native FICON interfaces. This can be done by point-to-point (direct connection), switched point-to-point (through a Fibre Channel Director), or cascaded FICON Directors (through two Fibre Channel Directors). FC channels can be defined as dedicated, MIF-shared, and spanned. See Figure 3-45 on page 97).

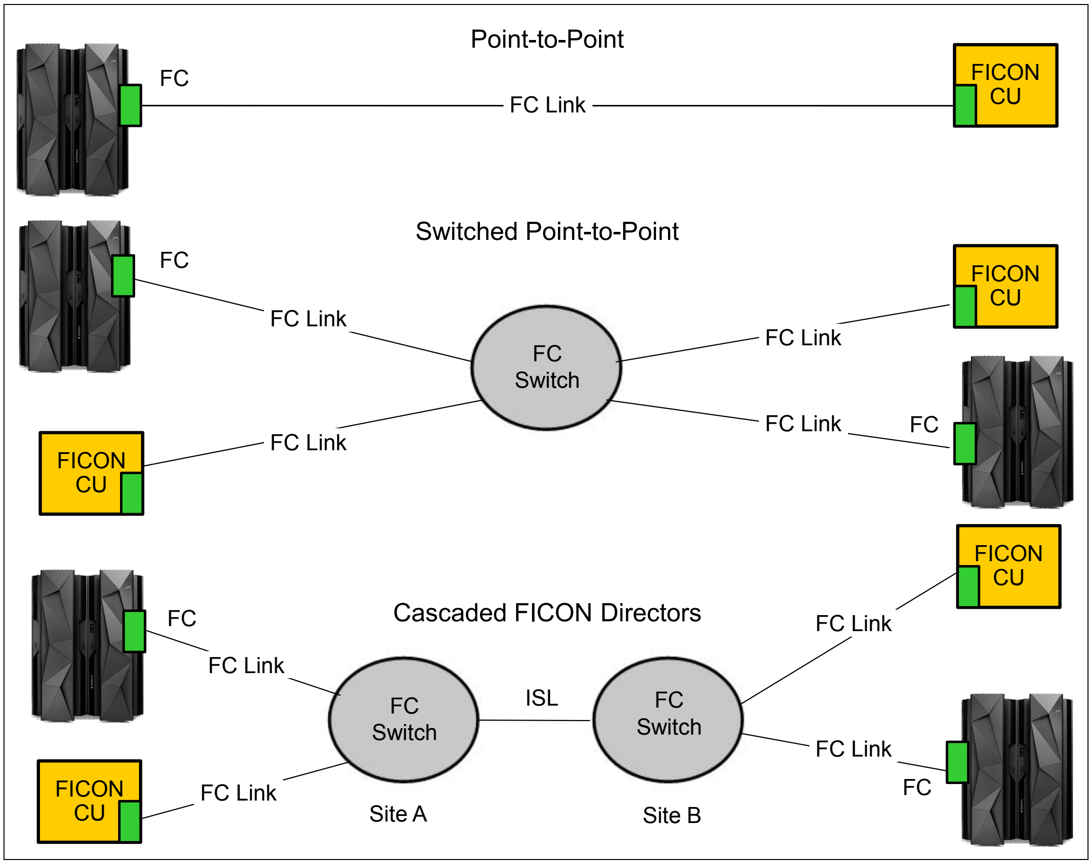

Figure 3-45 Supported FICON topologies

- Fibre Channel Protocol mode (FCP)

  This channel type allows for access to FCP devices over a single or multiple Fibre Channel switches to FCP devices. It also allows access over a single or multiple Fibre Channel switches to a Fibre Channel-to-SCSI bridge.

  The FICON Express16S and the FICON Express8S provide support for Fibre Channel and SCSI interfaces in Linux environments. FCP channels can be defined as dedicated, MIF-shared, and spanned. See Figure 3-46 on page 98.

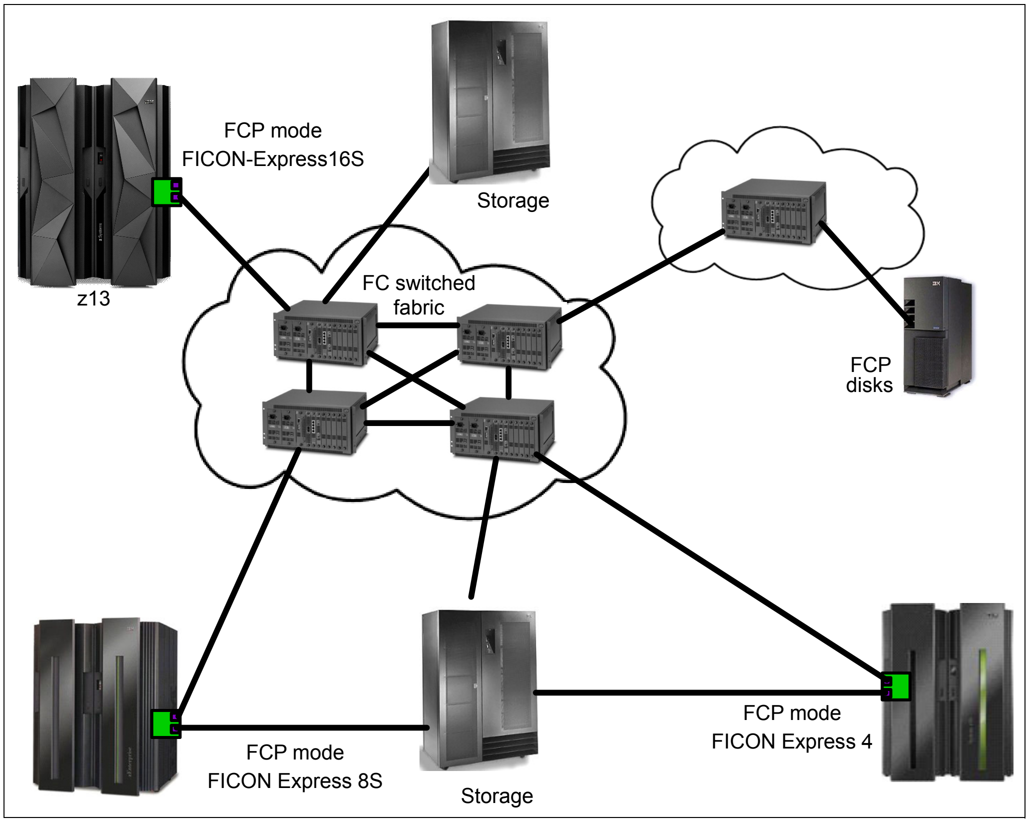

Figure 3-46 FCP connectivity topologies

3.3.4 Open System Adapter (OSA)

The z13 supports the OSA Express5S and OSA-Express4S features-carry forward.

The following OSA features can be ordered on a z13 server:

- OSA-Express5S 10GbE LR/SR
- OSA-Express5S GbE LX/SX
- OSA-Express5S 1000BASE-T

The following OSA features can be carried forward on an upgrade to z13:

- OSA-Express5S
- OSA-Express4S

The following features are *not* supported on a z13 server:

- OSA- Express3 GbE LX/SX
- OSA- Express3 1000BASE-T
- OSA- Express3 10 GbE LR/SR

A typical OSA network is shown in Figure 3-47 on page 99, which shows all the current supported OSA connectivity across z Systems servers.

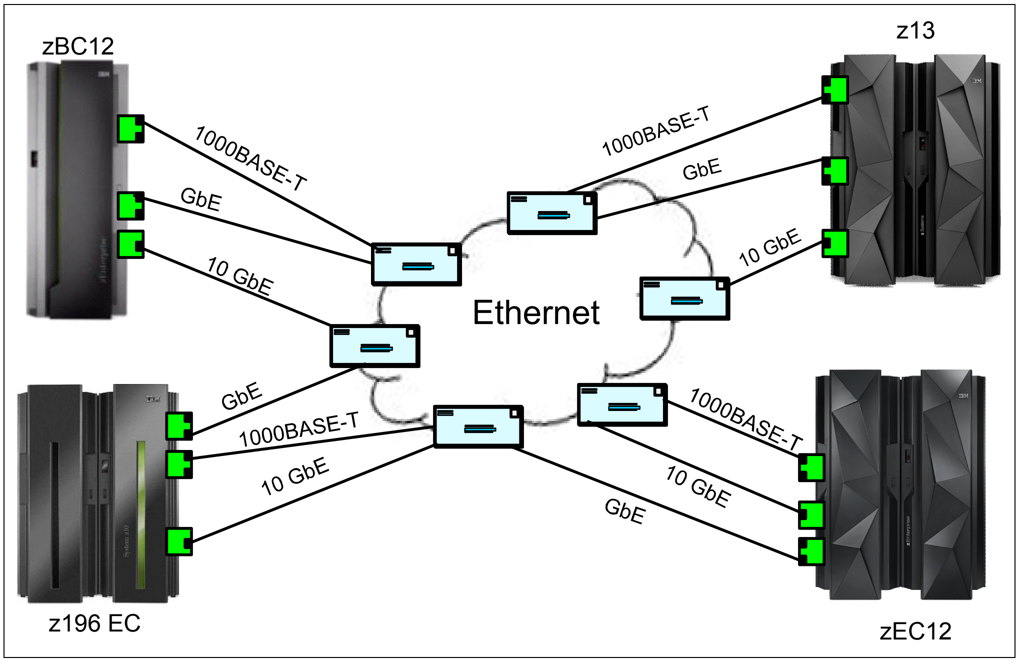

Figure 3-47 OSA supported connectivity

The OSA-Express5S 1000BASE-T Ethernet feature and the OSA-Express4S 1000BASE-T Ethernet feature can also be used for the OSA-Express Integrated Console Controller (OSA-ICC). The OSA-ICC provides a system console function when there is an IPL and operating systems support for multiple logical partitions. The OSA-ICC can be used as a replacement for any other console controller. The OSA-ICC is a standard feature on the 1000BASE-T Ethernet feature. To use it, code the channel type as OSC in the HCD.

The OSA-Express5S GbE, OSA-Express4S GbE and 1000BASE-T Ethernet features can provide direct (LPAR-to-LPAR) connectivity to the Communication Controller for Linux (CCL) on z196 and newer servers. They use the new Open Systems Adapter for Network Control Program (OSN) CHPID type. When the channel is defined as OSN (OSA-NCP), it can be used to replace IBM 3745 control units that run NCP. As with OSA-ICC, OSA-NCP is a standard feature. To use this feature for OSA-NCP, you must use a Linux logical partition that runs the appropriate application.

Channel type OSM, OSA for Unified Resource Manager requires an OSA-Express5S 1000BASE-T or an OSA-Express4S 1000BASE-T feature. Channel type OSM is used to build an INMN. The INMN provides control and management functions on a zBX and an z13 server to build a managed ensemble. It also provides a communication path for an HMC within an ensemble. For more information about managed ensembles, see 12.2.2, “Intranode management network (INMN)” on page 604.

Two ports on two different OSA-Express5S 1000BASE-T or OSA-Express4S 1000BASE-T adapters are used for redundancy. Only Port 0 on an OSA-Express5S 1000BASE-T or an OSA-Express4S 1000BASE-T card can be defined as channel type OSM. A copper CAT6 Ethernet cable is used to connect to the internal network in the z13 server.

Figure 3-48 on page 100 shows the port assignment of an OSA-Express5S 1000BASE-T adapter and the occupied ports if an OSM channel type is defined.

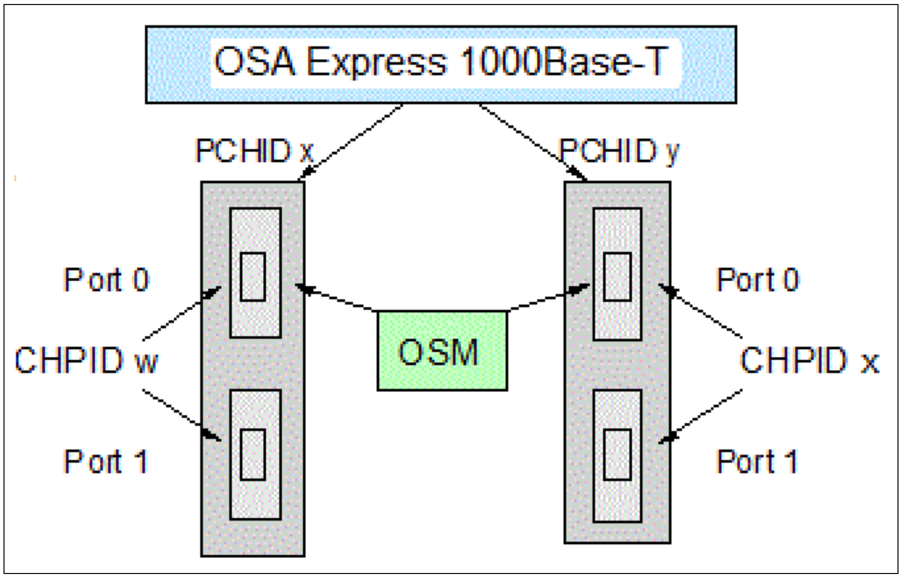

Figure 3-48 OSA-Express5S 1000BASE-T port assignment

For redundancy, two OSM CHPIDs are defined on two different OSA-Express5S 1000BASE-T or OSA-Express4S 1000BASE-T adapters. For each OSM CHPID, both ports (0 and 1) for the CHPID are used. However only port 0 on each adapter and CHPID is used to connect to the network.

To build a managed ensemble, an extra channel type OSX must be defined on an OSA-Express5S 10 GbE or OSA-Express4S 10 GbE feature-carry forward. The OSX channel type is used to build an IEDN. The IEDN allows communication and data sharing between z/OS images and provides a communication path to a zBX. For redundancy, two OSA-Express5S 10 GbE or OSA-Express4S 10 GbE carry-forward features are required. The OSX channel type connects to the Top of Rack (TOR) switch in a zBX.

Figure 3-49 shows the port assignment of an OSA-Express5S 10 GbE adapter and the occupied ports if an OSX channel type is defined.

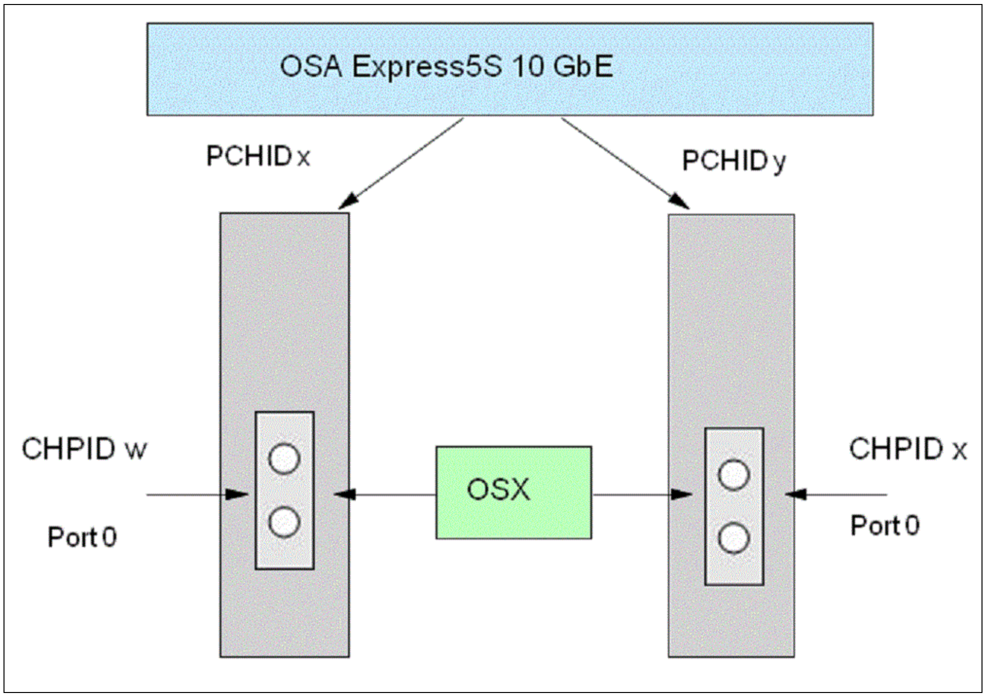

Figure 3-49 OSA-Express5S 10 GbE port assignment

For redundancy, two OSX CHPIDs are defined on two different OSA-Express5S 10 GbE or OSA-Express4S 10 GbE adapters. For each OSX CHPID, one port for the CHPID is occupied and used to connect to the zBX. For more information, see 12.2.3, “Intraensemble data network (IEDN)” on page 609.

3.3.5 HiperSockets

HiperSockets provide the fastest TCP/IP communications between z/OS, z/VM, z/VSE, and Linux logical partitions within an z13. HiperSockets provide internal “virtual” local area networks that act like IP networks inside a z13 server. This effect is achieved by using the Licensed Internal Code (LIC) and supporting device drivers on the operating systems. HiperSockets establishes 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 hence 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, bypassing external network delays and provides users high-speed logical LANs with minimal system and network overhead.

HiperSockets can be defined as MIF-shared in a CSS and as spanned channels across multiple CSSs. A HiperSockets CHPID can be seen as an “internal LAN” to the server. The level of sharing is determined by the 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 Systems

On a z13, 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 are using 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 both 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 resulting in increasing performance and a simplified the 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 z13 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 has been increased by allowing streaming of bulk data over a HiperSockets link between logical partitions. Multiple writes with fewer I/O interrupts reduces processor usage of 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, 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

- HiperSockets integration with the IEDN

  The z13 and the zEnterprise systems provide the capability to integrate HiperSockets connectivity to the IEDN. This configuration extends the reach of the HiperSockets network outside the CPC to the entire ensemble and is displayed as a single Layer 2. Because HiperSockets and IEDN are both internal z Systems networks, the combination allows z Systems virtual servers to use the optimal path for communications.

- zIIP-Assisted HiperSockets for large messages

  In z/OS, HiperSockets has been 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 Systems Connectivity Handbook*, SG24-5444.

3.3.6 Integrated Coupling Adapter (PCIe-O SR) Adapter

The z13 servers now offers the benefit from the high speed and low latency that are offered by Integrated Coupling Adapter (ICA) PCIe-O SR technology. It provides improved reliability, scalability, and performance which is important for both Parallel Sysplex and I/O interconnectivity.

Integrated Coupling Adapter (ICA) is a point-to-point interconnect architecture that was developed to satisfy contemporary requirements for higher bandwidth and ability to scale with increasing bandwidth demand. Integrated Coupling Adapter is a PCIe-Gen3 fanout card in the CPC drawer. It is preferred for short distance coupling z13 to a z13. It is not available on older servers. There is no performance degradation compared to Coupling over InfiniBand 12X IFB3 protocol.

PCIe-Gen3 fanout card details:

- Short reach adapter, distance up to 150 m
- Up to 32 ports maximum
- IOCP Channel Type = CS5
- Feature code 0172, 2 ports per adapter
- Up to 4 CHPIDs per port, 8 per feature, 8 buffers (that is 8 subchannels) per CHPID
- ICA requires new cabling for single MTP connector which differs from 12X InfiniBand split transmit/receive connector

Software requirements are z/OS V2.1, V1.13 or V1.12 with PTFs for APARs OA44440 and OA44287

The advantage to using ICA connectivity is that it provides more ICA coupling fanouts per CPC drawer when compared to 12X Coupling over InfiniBand on either z196, zEC12 or z13. A single z13 CPC drawer supports up to 20 ICA links versus 16 12X on z196/zCE12, 8 12X on z13. Utilizing ICA frees HCA slots for essential Coupling over InfiniBand during migration.

For z13 to z13 connectivity using ICA in place of Coupling over InfiniBand might enable clients to remain in the same CPC footprint as their z196 or zEC12 enterprises.

3.3.7 InfiniBand (PSIFB) Adapter

The z13 servers continue to benefit from the high speed and low latency that are offered by InfiniBand technology. It provides improved reliability, scalability, and performance, which is important for both Parallel Sysplex and I/O interconnectivity.

InfiniBand is a point-to-point interconnect architecture that was developed to satisfy contemporary requirements for higher bandwidth and ability to scale with increasing bandwidth demand. The InfiniBand specification defines the raw bandwidth of one lane (referred to as 1x) connection at 2.5 Gbps. Two more lane widths are specified, referred to as 4x[3] and 12x, as multipliers of the base link width.

Like Fibre Channel, PCI Express, Serial ATA, and many other contemporary interconnects, InfiniBand is a point-to-point, bidirectional serial link. It is intended for the connection of processors with high-speed peripheral devices, such as disks. InfiniBand supports various signaling rates and, as with PCI Express, links can be bonded together for more bandwidth.

The serial connection's signaling rate is 2.5 Gbps on one lane in each direction, per physical connection. Currently, InfiniBand also supports 5 Gbps or 10 Gbps signaling rates.

InfiniBand links use 8b/10b encoding (every ten bits sent carry eight bits of data). Therefore, the useful data transmission rate is four-fifths of the signaling rate (signaling rate equals raw bit rate). Thus, links carry 2, 4, or 8 Gbps of useful data for a 1x link.

InfiniBand links can be aggregated in units of 4 or 12, indicated as 4x[3] or 12x. A 12x link therefore carries 120 Gbps raw or 96 Gbps of payload (useful) data. Larger systems with 12x links are typically used for cluster and supercomputer interconnects, as implemented on the z13, and for inter-switch connections.

For more information and the standard for InfiniBand, see the InfiniBand website:

http://www.infinibandta.org

Exception: Not all properties and functions that are offered by InfiniBand are implemented on the z13. Only a subset is used to fulfill the interconnect requirements that were defined for z13.

InfiniBand currently defines several link speeds at the physical layers:

- Single Data Rate (SDR), delivering 2.5 Gbps per physical lane
- Double Data Rate (DDR), delivering 5.0 Gbps per physical lane
- Quadruple Data Rate (QDR), delivering 10.0 Gbps per physical lane

Combining link speeds with interface widths gives the link or signaling rates that are listed in Table 3-8.

Table 3-8 Interface width and link ratings

Width	Single Data Rate	Double Data Rate	Quadruple Data Rate
1X	2.5 Gbps	5.0 Gbps	10 Gbps (1 GBps)
12X	30.0 Gbps (3 GBps)	60.0 Gbps (6 GBps)	120 Gbps (12 GBps)

[3] z13 does not support this data rate.

InfiniBand coupling links

InfiniBand (IFB) coupling links are high-speed links on z13, zEC12, z196 and z114. The IFB coupling links originate from these types of fanouts:

- HCA3-O (FC 0171) for 12x IFB links on zEnterprise CPCs
- HCA3-O LR (FC 0170) for 1x IFB links on zEnterprise CPCs

PSIFB links are supported from z/OS V1.12 and later. The information necessary to identify the required service is available in the Preventive Service Planning (PSP) buckets named 2817DEVICE, 2827DEVICE, and 2964DEVICE for z196, zEC12 and z13 respectively. Review the PSP information thoroughly early in the planning process to allow time for ordering and installing any necessary software.

IFB

z13 supports one type of 12x InfiniBand coupling link: FC 0171 HCA3-O (12xIFB) fanout.

The HCA3-O fanout support IFB coupling link operates at 6 GBps (12x IFB). IFB coupling links use a fiber optic cable that is connected to a HCA3-O fanout. The maximum distance for an IFB link is 150 meters. The fiber cables are industry standard OM3 50/125 micrometer multimode optical cables with Multifiber Push-On (MPO) connectors. IFB coupling links are defined as CHPID type CIB. Up to 16 CHPIDs across two ports on one HCA-O fanout.

12x IFB3 protocols

The HCA3-O feature that supports 12x InfiniBand coupling links is designed to deliver improved services times. When you use the 12x IFB3 protocol, synchronous service times are 40% faster than when you use the 12x IFB protocol.

IFB LR

z13 supports one type of 1x InfiniBand coupling link: FC 0170 HCA3-O LR (1xIFB) fanout.

Each connection supports a link rate of 5 Gbps if connected to a z13, zEC12, z196, or a z114 server, or to a z Systems qualified Dense Wave Division Multiplexing (DWDM). The connections support a data link rate of 2.5 Gbps when connected to a z Systems qualified DWDM. The link rate is auto-negotiated to the highest common rate. 1x IFB coupling links use a fiber optic cable that is connected to a HCA3-O LR fanout. The maximum unrepeated distance is 10 km, and up to 100 km with z Systems qualified DWDM. The fiber cables that are used for 1x IFB links are standard 9-µm single mode fiber optic cables with an LC Duplex connector.

The HCA3-O LR fanout has four ports for optical link connections. It supports up to 16 CHPIDs across all ports on a fanout. IFB LR coupling links are defined as CHPID type CIB, supporting 7 or 32 subchannels per CHPID.

> **Note:** All types of Integrated Coupling Adapter (PCIe-O SR) and PSIFB coupling links support STP.

Host Channel Adapter (HCA2-C) (Carry Forward Only)

z13 supports the connectivity to the I/O drawer when carried over from z196 or zEC12. The connectivity to the I/O drawers (I/O domains) in the z13 server is supported by InfiniBand technology. This connection provides a data rate of 6 Gbps.

An HCA2-C fanout in the front of a CPC drawer connects over a 12x InfiniBand copper cable of 1.5 - 3.5 meters long to the IFB-MP card in the I/O drawer. There is a passive connection in the IFB-MP to provide the redundancy for the I/O interface. This setup allows for concurrent repairs against the cabling, the HCA-2C fanout, or the CPC drawer.

The use of InfiniBand for the I/O interconnectivity is entirely internal to the z13 server and requires no further planning or actions.

Note: I/O cages cannot be carried forward to z13.

3.3.8 Fanouts

The z13 server uses fanout cards to connect the I/O hardware subsystem to the processor drawers. They also provide the Integrated Coupling Adapter and InfiniBand coupling links for Parallel Sysplex. All fanout cards support concurrent add, delete and move.

Internal z13 I/O infrastructures are established by PCIe fanout cards and InfiniBand fanout cards:

- The PCIe Generation3 fanout card is a one port card and connects to a PCIe I/O drawer supporting an eight slot I/O domain. This one port card is always installed in pairs in order to support an I/O domain.
- The InfiniBand HCA2-C fanout card is available for carry forward of FICON Express8 features.

The PCIe and InfiniBand fanouts are on the front of each processor drawer. Each processor drawer has ten PCIe Generation3 slots and four InfiniBand fanout slots. The PCIe slots are named LG02 through LG06 and LG11 through LG15, left to right. The InfiniBand fanout slots reside in the middle of the processor drawer and are named LG07 through LG10, left to right. Slots LG01 and LG16 are not used for I/O infrastructure. Five types of fanout cards are supported by z13. Each slot can hold one of the following five fanouts:

- PCIe Generation3 fanout card: This copper fanout provides connectivity to the PCIe switch card in the PCIe I/O drawer.
- Host Channel Adapter (HCA2-C): This copper fanout provides connectivity to the IFB-MP card in the I/O drawer. This is a carry forward only option.
- Integrated Coupling Adapter (ICA): This adapter provides coupling connectivity between two z13 servers.
- Host Channel Adapter (HCA3-O (12xIFB)): This optical fanout provides 12x InfiniBand coupling link connectivity up to 150-meter (492-ft) distance to a z13, zEC12, zBC12, z196, and z114.
- Host Channel Adapter (HCA3-O LR (1xIFB)): This optical long range fanout provides 1x InfiniBand coupling link connectivity up to a 10 km (6.2 miles) unrepeated (or 100 km (62 miles) when extended by using z Systems qualified DWDM equipment) distance to z13, zEC12, zBC12, z196 and z114.

The PCIe Generation3 fanout card includes one port; the HCA3-O LR (1xIFB) fanout includes four ports; other fanouts include two ports.

Figure 3-50 on page 107 illustrates I/O interconnectivity for PCIe I/O drawers and IO drawers.

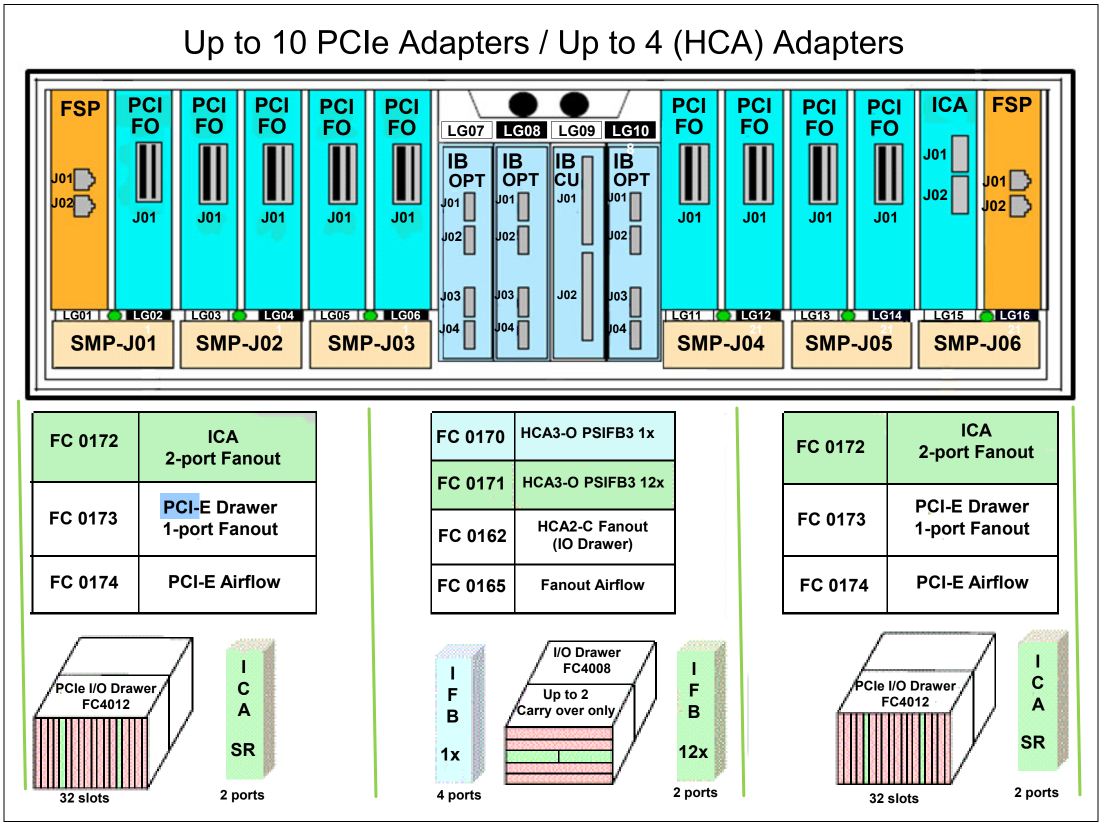

Figure 3-50 PCIe and InfiniBand I/O infrastructure

Fanout considerations

Fanout slots in each processor drawer can be used to plug different fanouts, where each fanout is designed for a special purpose. One drawer can hold up to ten PCIe fanouts and four InfiniBand fanout cards.

For migration purposes, the number of available HCA3 or HCA3 LR cards in the z13 can be very important (for coupling connectivity to zEC12, zBC12, z196, and z114). The carry forward of FICON Express8 features reduces the number of available InfiniBand fanout slots by two. This limitation can be important for Coupling connectivity planning. There are only two fanout adapter slots for HCA3 or HCA3 LR features available. For a three or four processor drawer the IFB fanout capability is relieved because the total number of IFB fanout slots increases with each additional processor drawer.

3.3.9 Coupling links

The z13 server supports two types of external coupling links: Integrated Coupling Adapter (PCIe-O SR) and PSIFB coupling links:

- Integrated Coupling Adapter (PCIe-O SR)
  - Maximum distance up to 150 m
  - Up to 32 ports maximum
  - OM4 (24 fiber, 4700 MHz-km 50/125 micron fiber with MTP connectors-150 meters OM3 (24 fiber, 2000 MHz-km 50/125 micron fiber with MTP connectors-100 meters
  - Channel Type = CS5
  - Requires new cabling for single MTP connector (differs from 12X InfiniBand split Transmit/Receive connector)

- PSIFB
  - HCA3-O fanout
    - Maximum distance of 150 m (492 feet)
    - 50 micron OM3 (2000 MHz-km) multimode fiber cable
    - MPO connector
    - Channel Type = CIB
    - Maximum of 32 links per server
    - Link data rate 6 GBps
  - HCA3-O LR fanout
    - Maximum distance of 10 km (6.2 miles) or 100 km (62 miles) with DWDM
    - 9 micron single-mode (SM) fiber cable
    - LC duplex connector
    - Channel Type = CIB
    - Maximum of 64 links per server
    - Link data rate 5 GBps or 2.5 Gbps with DWDM

> **Notes:** ISC-3 coupling links are not supported on the z13. The zEC12 was the last platform to offer ISC-3 coupling links.

The following channel types can be defined for coupling links:

- PCIe-O SR (CS5): This CHPID type is used to identify a coupling over Integrated Coupling Adapter channel path. The CHPID can be DED, REC, SHR, or SPAN. The PCIe-O SR link supports point-to-point up to 150 m.
- PSIFB (CIB): This CHPID type is used to identify a Coupling over InfiniBand (CIB) channel path. The CHPID can be DED, REC, SHR, or SPAN. The HCA3-O link supports point-to-point up to 150 m. The HCA3-O LR link supports point-to-point up to 10 km.
- Internal Coupling links in peer mode (ICP): Internal Coupling channels are linkless connections (controlled by Licensed Internal Code) that do not need any cabling. They provide connectivity between a Coupling Facility and logical partitions that run on the same CPC.

The following are high-level configuration guidelines for CS5 CHPIDs:

- No PCHID can be specified for a CS5 channel path; instead an AID and PORT are used.
- A spanned CS5 CHPID requires the same AID and PORT to be specified for all channel subsystems where it is defined.
- Up to 8 CS5 CHPIDs can be defined to the same AID (4 CHPIDs per port). This configuration is verified at the HCD level. You cannot define more than four CHPIDs per port.
- All CS5 CHPIDs must be connected within HCD/HCM before a production IODF can be built.
- A CS5 CHPID can be connected only to another CS5 CHPID.

> **Note:** A configuration that connects two CS5 CHPIDs within the same system using the same AID and port is not supported.

The following high-level configuration guidelines are for CIB CHPIDs:

- No PCHID can be specified for a CIB channel path. An AID and PORT are used.
- A spanned CIB CHPID requires the same AID and PORT to be specified for all channel subsystems where it is defined.

- Up to 16 CIB CHPIDs can be defined to the same AID. This configuration is verified at the HCD level, meaning all 16 can be defined on one port. For optimum performance, especially when you use HCA3-O links, do not define more than four CHPIDs per port. However, if the link is being used to provide connectivity for sysplexes with low levels of Coupling Facility activity, you can assign more than four CHPIDs per port.
- A combination of CIB and ICP CHPIDs are allowed on the same control unit (CF image), up to the maximum of eight paths per CU.
- All CIB CHPIDs on a single control unit must have the same connecting system (CSYSTEM).
- A CIB CHPID can be connected only to another CIB CHPID.
- All CIB CHPIDs must be connected within HCD/HCM before a production IODF can be built.
- A processor that has a connected CIB channel path must have an established local system name (LSYSTEM). This is handled by HCD.

> **Note:** A configuration that connects two CIB CHPIDs within the same system using the same AID and port is not supported. Although prior versions of IOCP/HCD allow this type of configuration, the current version of IOCP/HCD prevents such a definition.

All the supported coupling link types can coexist on a single server and between two systems.

For z13, the maximum number of coupling CHPIDs is 256 and the maximum number of Coupling Facility Image CHPIDs remains at 128.

Support for multiple CHPIDs per port

With PCIe-O SR and PSIFB coupling technology, you can define multiple CHPIDs on a single port. HCA3-O fanout has two ports and HCA3-O LR fanout has four ports. You can define 16 CHPIDs on one card in any combination

PCIe-O SR fanout has two ports and you can define a maximum of 4 CHPIDs on each port.

Use this new support for the following reasons:

- PCIe-O SR and PSIFB links can be shared by multiple Sysplexes. For PSIFB links, multiple CHPIDs can share a physical link (up to 16 per adapter). For PCIe-O SR links, multiple CHPIDs can share a physical link (up to 8 per adapter).
- Defining more coupling link CHPIDs provides more subchannels for CF communication. Each peer-mode coupling link CHPID defines seven subchannels that support seven concurrent I/O operations to the coupling facility. If the volume and duration of CF accesses are high enough to cause subchannel busy conditions, defining more CHPIDs can improve CF communication performance. This action might also provide significant relief when multiple z/OS images on a single system share a coupling link CHPID over a MIF. For heavily used systems, this is often a source of subchannel busy conditions.

> **Consideration:** Defining multiple CHPIDs on a single physical coupling link does not satisfy the condition of having multiple coupling links between coupled systems. Configure multiple physical links, on different HCAs or ICAs, to avoid single points of failure.

Coupling link maximum values

Both ICA and PSIFB links connect a port on an HCA over a physical cable to the equivalent one on a different server. The z13 supports up to 128 coupling link CHPIDs. The maximum number of coupling links that are allowed for each type is summarized in Table 3-9.

Table 3-9 Coupling link server support

Link type		Maximum supported links			
	z13	zEC12	z196	z114	
				M10	M05
IC	32	32	32	32	32
ICA	32	N/A	N/A	N/A	N/A
HCA3-O LR (1x IFB)	64	64	48	32	16
HCA3-O (12x IFB)	32	32	32	16	8

The maximum number of coupling links in a particular server is affected by many factors including the number of I/O drawers and CPC drawers installed. PSIFB links (HCA3-O 1x IFB and HCA3-O 12x IFB) contend for adapter space. Total port counts vary depending upon mix of 1x and 12x links configured and will never exceed the single 1x maximum of 64 ports. PSIFB and ICA links type do not contend with each other for adapter space. The maximum number of external links that are supported on z13, PSIFB and PCIe-O SR combined, is 96.

No combined limit on physical coupling links exists, but the maximum number of coupling CHPIDs defined is 128 per LPAR. Defining IC links, and defining multiple CHPIDs over the installed PSIFB and PCIe-O SR ports might reach the limit on the number of CHPIDs before using all the physical links. Even though these maximums have increased from 128 to 256 on z13, consider them when you are planning to use the flexibility allowed by the PSIFB links.

CIB CHPIDs can be shared or spanned across logical partitions and channel subsystems on the z13, zEC12, and z196. CS5 CHPIDs can be shared or spanned across logical partitions and channel subsystems on a z13.

Coexistence

When you implement PCIe-O SR or PSIFB, the requirements for z13 server coexistence must be considered. The z13 can coexist *only* with zEnterprise zEC12 and z196 servers in a Parallel Sysplex. Any earlier servers (such as z800, z900, z890, z990, z9 and z10) in the Parallel Sysplex environment must be removed or replaced by a supported server before the z13 can be added.

PCIe-O SR coupling links are available *only* on z13. PSIFB coupling links are available on zEC12 and z196 servers. Table 3-10 on page 111 lists the supported coupling link options for connecting a z13 server to another z13, zEC12 and z196 server.

Table 3-10 Coupling link options to connect z13 to a z13, zEC12 and z196

Connect to	**z13**		
	PCIe-O SR	**PSIFB 12x IB DDR**	**PSIFB LR 1x IB SDR**
z13	Yes	Yes	Yes
zEC12	No	Yes	Yes
z196	No	Yes	Yes

The coupling link configuration options are as follows:

- z13 to z13
- z13 to zEC12
- z13 to z196 and z114

On a z13 server, up to 96 physical external coupling links can be ordered. A maximum of 128 CHPIDs can be configured across the physical links.

Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors that include latency through the adapters, cable lengths, and the type of workload.

When you configure a z13 in a Parallel Sysplex configuration, ensure that the other servers have the correct channel type and definition to connect to a z13.

Adapter ID (AID) and port assignment

Unlike channels that are installed in an I/O drawer which are identified by a PCHID number related to their physical location, PCIe-O SR and PSIFB fanouts and ports are identified by an AID. The AID is used to define CHPIDs for PCIe-O SR and PSIFB links just as PCHIDs are used to define CHPIDs for other types of definitions. The assignment of AIDs is covered in “Adapter ID (AID) assignment” on page 112.

The AID assigned to an adapter can be found in the *PCHID report* that is provided for a new server or for MES upgrades on existing servers. This is the same report that provides the PCHID information for other types of definitions. Example 3-1 on page 112 shows part of a PCHID report for a z13 model N63. In this example, eight adapters are installed: four in the first CPC drawer (location A19) and the other four in the second CPC drawer (location A15). The adapters are installed in individual slots in each drawer. Note the four HCA3-O 12X features (FC0171), two HCA3-O 1X features (FC0170), and two PCIe-O SR features (FC0172) in this CPC.

Example 3-1 Sample PCHID REPORT showing AID assignments

```
CHPIDSTART
  21751417                      PCHID REPORT                        Aug 05,2014
  Machine: 2964-N63  SN1
  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
  Source          Cage  Slot  F/C    PCHID/Ports or AID               Comment
  A19/LG07        A19A  LG07  0171   AID=08
  A19/LG09        A19A  LG09  0171   AID=0A
  A19/LG10        A19A  LG10  0170   AID=0B
  A15/LG07        A15A  LG07  0171   AID=0C
  A15/LG09        A15A  LG09  0171   AID=0E
  A15/LG10        A15A  LG10  0170   AID=0F
  A19/LG05        A19A  LG05  0172   AID=27
  A15/LG12        A15A  LG12  0172   AID=34
```

Remember: When a PCIe-O SR or PSIFB HCA fanout card is moved from one slot location on a processor book to another, the AID moves also (that is, the AID is *retained*).

LSYSTEM, CSYSTEM, and CPATH

A local system name (LSYSTEM in the IOCP) is required on every server that uses PCIe-O SR and PSIFB links. The LSYSTEM unique keyword is on the ID statement. The local system name is used by the CPC to identify itself when it establishes a coupling facility connection using CIB and CS5 channel paths. Generally, use the same CPC name that is specified for the processor on the Hardware Management Console. This action helps avoid naming conflicts.

CSYSTEM is on the CHPID statement. This is the name given to the attached system by the LSYSTEM value in the IOCP of the system with which this CHPID is going to connect.

The CPATH is also defined on the CHPID statement. The value of this keyword, similar to ICP, identifies the CSS and CHPID on the connecting system (CSYSTEM) to which this channel path is to communicate. The CPATH keyword is valid only for ICP, CIB and CS5 channel paths, and is required for all ICP, CIB and CS5 definitions. An ICP, CIB or CS5 channel path must connect to another channel path of the same type. It can connect to only one path. A channel path cannot connect to itself. The CPATH keyword has extra syntax rules when the connected channel paths belong to the same CPC (that is, they are internal connections). Channel paths cannot connect to each other if they both have candidate lists with the same, single logical partition. This prevents the definition of ICP, CIB and CS5 channels in a configuration with only one logical partition defined.

Adapter ID (AID) assignment

The Adapter ID (AID) is a number that is assigned to every PCIe and PSIFB fanout at installation time. It is unique for the system. There is only one AID assigned per fanout so all the ports on the fanout share that AID. In the IOCP, the AID must be used to connect the physical location of the fanout with the logical assigned CHPID. For assignment and handling, the AID is bound to the serial number of the fanout. If the fanout is moved, the AID moves with it so it is not necessary to readjust the IOCDS. The assignments of the AIDs are found in the PCHID Report for new servers (or MES upgrades) and in the HMC and SE windows after installation.

For new z13 system builds and newly installed books, the AID can be determined by using the data in Table 3-11.

Table 3-11 Initial AID number assignments

CPC drawer	Location	Fanout slot	AID
First	A15A	LG02-LG06 (PCIe)	2E-32
		LG07-LG10 (IFB)	0C-0F
		LG11-LG15 (PCIe)	33-37
Second	A19A	LG02-LG06 (PCIe)	24-28
		LG07-LG10 (IFB)	08-0B
		LG11-LG15 (PCIe)	29-2D
Third	A23A	LG02-LG06 (PCIe)	1A-1E
		LG07-LG10 (IFB)	04-07
		LG11-LG15 (PCIe)	1F-23
Fourth	A27A	LG02-LG06 (PCIe)	10-14
		LG07-LG10 (IFB)	00-03
		LG11-LG15 (PCIe)	15-19

STP links

Integrated Coupling Adapter (PCIe-O SR) and PSIFB links can be used to pass time synchronization signals by using STP. The same coupling links can be used to exchange timekeeping information and Coupling Facility messages in a Parallel Sysplex.

STP provides a coordinated time source for systems that are connected over coupling links. It replaces the Sysplex Timer as a time source for interconnected systems. STP uses coupling links to transmit time signals between interconnected systems. Integrated Coupling Adapter (PCIe-O SR) and PSIFB links are supported for STP timing signals between systems that support Integrated Coupling Adapter (PCIe-O SR) and PSIFB, including timing-only links where required. Coexistence with other types of coupling links that are used for STP is also supported. For more information, see the *Server Time Protocol Planning Guide*, SG24-7280.

Consideration: Use at least two physical coupling links between any two servers that exchange STP messages to prevent a single point of failure.

Coupling links are designed to send a reliable unambiguous “going away signal.” This signal indicates that the server on which the coupling link is running is about to enter a failed (check stopped) state. When the “going away signal” sent by the Current Time Server (CTS) in an STP-only CTN is received by the Backup Time Server (BTS), the BTS can safely take over as the CTS. It does not have to rely on the previous offline signal (OLS) in a two-server CTN or the Arbiter in a CTN with three or more servers.

This enhancement is exclusive to zEnterprise CPCs and z13. It is available only if you have an PCIe-O SR or HCA3-O (12xIFB) or HCA3-O LR (1xIFB) on the CTS communicating with an PCIe-O SR or HCA3-O (12xIFB) or HCA3-O LR (1xIFB) on the BTS. The already available STP recovery design is still available for the cases when a going away signal is not received or for other failures besides a server failure.

Coupling Facility Control Code (CFCC)

The new Coupling Facility Control Code (CFCC) level 20 is introduced by the z13 server. CFCC level 20 allows allocation of up to 2047 structures. CFCC level 20 can coexist with previous CFCC levels on zEC12 and z196 servers. The minimum CFCC level is 19 for zEC12 and CFCC level 17 for z196.

To calculate structure sizes, use the Coupling Facility Structure Sizer (CFSizer) tool. The CFSizer returns structure sizes that are based on CFCC level. The CFSizer tool is available at the following web page:

http://www.ibm.com/systems/z/cfsizer

3.3.10 Configuration rules

Table 3-12 summarizes the maximum number of CHPID types supported on a z13 server.

Table 3-12 z13 maximum number of supported CHPIDs

Maximum channel paths (CHPID type)	z13
Channel path total	1536
FC and FCP (assuming no FICON Express 8)	320
OSC, OSD, OSE, OSN, OSM[a], and OSX	96
CS5, ICP, and CIB combined	128 or 256
CS5	128
ICP	32
CIB	128 or 256
IQD	32

a. Only two OSM CHPIDs are required to configure an ensemble.

The following configuration rules are implemented and enforced by using the HCD and IOCP. All control units and I/O devices that attach to the server must be defined to the channel subsystem. When you define the I/O configuration for the server CSS over HCD/IOCP, you must specify these items:

- For logical partitions, the logical partition name, CSS ID, and MIF ID
- Channel paths and their assignment to CSSs and logical partitions
- FICON Director, where appropriate
- Control units that are attached to the channel paths
- I/O devices that are assigned to the control units

Table 3-13 lists the z Systems platform attributes.

Table 3-13 z Systems server attributes

Constant machine attributes		z13, zEC12, z196 and z114
Maximum configurable physical control units (PCUs)	PCUs per CVC[a], CBY[a]	48
	PCUs per OSD	16
	PCUs per OSE, OSC, OSN	1
	PCUs per OSM, OSX	16
	PCUs per CFP, CBP[b], ICP, CIB, CS5	1
	PCUs or link addresses per CNC[a], CTC	120
	PCUs or link addresses per FC	256
	PCUs per FCP	1
	CUs per IQD	64
Maximum configurable devices	Per CFP, ICP, CIB	7
	Per CS5	8
	CIB (1x IFB for zEnterprise CPCs only)	32
	Per CNC	1024
	Per CTC	512
	Per OSC[c]	253
	Per OSD	1920
	Per OSE	254
	Per OSM, OSX	1920
	Per OSN	480
	Per FCP	480
	Per FC (FICON Express 16S)	32K
	Per FC (FICON Express 8S)	24K
	For all IQD channel paths	12K

a. Not supported on z13 and zEC12.
b. Not supported on z13, zEC12, z196, and z114.
c. A limit of 120 sessions can be defined at the HMC/SE.

4

I/O Autoconfiguration function (zDAC)

In this chapter, we describe how to use the I/O Autoconfiguration function with the IBM z13 and z/OS Version 2 Release 1.

I/O Autoconfiguration is also known as *z/OS Discovery and Autoconfiguration (zDAC)*.

With the I/O Autoconfiguration function, Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) allow you to discover undefined FICON storage devices on a direct access storage device (DASD) or an enterprise tape solution connected to the platform either direct attached, or using a switch.

You can define policies before using I/O Autoconfiguration, or use the defaults provided with HCD/HCM. Depending on these policies, I/O Autoconfiguration function can automatically define a control units devices and include them into a specified target work input/output definition file (IODF), or propose definitions. The user can accept the proposed definition without changes, or update the proposed definitions before committing them to the specified target work IODF.

HCD and HCM provide these capabilities:

- A dialog to define auto configuration policies
- A dialog to run the discovery and definition process

The I/O Autoconfiguration function is available on z196, zEC12, and z13.

This chapter includes the following topics:

- Description of I/O Autoconfiguration
- Enhancements introduced with z/OS Version 2 Release 1
- Prerequisites and operational considerations for using I/O Autoconfiguration
- Overview of the I/O Autoconfiguration process
- I/O Autoconfiguration using HCD
- I/O Autoconfiguration using HCM

4.1 Description of I/O Autoconfiguration

You can use the Hardware Configuration Definition (HCD) I/O Autoconfiguration function to run automatic configuration tasks to define switched FICON connected DASD and tape control units and devices. These must not be defined in either the active or currently accessed IODF. HCD can start the I/O subsystem to discover I/O hardware in the current configuration that is accessible to the system. Proposed definitions are automatically written into a specified target work IODF that is created as a copy of the active or accessed IODF.

During I/O Autoconfiguration processing, HCD presents the discovered controllers, control units, and devices to the user and proposes way to configure them. You can accept or change these definition proposals. The configuration definitions are then written into the specified target IODF.

HCD provides a series of panels to run the automatic I/O configuration. These discovery panels are scoped by policies, and the parameters for these policies are accessed with HCD option 0 or HCM. You must have an HCD profile data set (HCDPROF) allocated to your HCD session or HCM session for these policy parameters to take effect and retained for future discoveries.

I/O Autoconfiguration is available starting with the z196. It requires the same access authorization as used for dynamic IODF activation.

4.2 Enhancements introduced with z/OS Version 2 Release 1

HCD and HCM (HCD/HCM) provide the I/O Autoconfiguration function in z/OS. In addition to switched FICON connected controllers, I/O Autoconfiguration in z/OS V2R1 can discover FICON directly attached controllers and devices, and proposes point-to-point connection paths if available.

I/O Autoconfiguration supports the inclusion or exclusion of specific switches or CHPIDs into the discovery and proposal process, which users can explicitly specify with the invocation of an I/O Autoconfiguration request. For this purpose, HCD/HCM introduce four autoconfiguration keywords.

- AUTO_CHPID_INCLUDE
- AUTO_CHPID_EXCLUDE
- AUTO_SWAD_INCLUDE
- AUTO_SWAD_EXCLUDE

The Autoconfiguration policy keyword AUTO_SS_DEVNUM_SCHEME accepts a new value NONE. This value bypasses control unit and device number proposals by HCD/HCM and lets the user manually apply the numbers for detected objects.

I/O Autoconfiguration allows discovery by controller serial number and filters the discovered controllers accordingly. HCD/HCM can process an I/O Autoconfiguration request that is partially directed against unavailable systems of an LPAR group or a sysplex, or against systems that are not capable of supporting I/O Autoconfiguration. Users can specify that the request applies to appropriate systems only, and that unavailable/incapable systems are tolerated but ignored. HCD/HCM allows users to change certain I/O Autoconfiguration policies between two subsequent controller discoveries without the need to make a new fabric discovery. This enables I/O Autoconfiguration to perform each new controller discovery with changed policies.

4.3 Prerequisites and operational considerations for using I/O Autoconfiguration

Consider the following information when you use the I/O Autoconfiguration function of HCD:

- The I/O Autoconfiguration process minimally requires that the LPAR (LP) groups are running on z196, zEC12, or z13 platforms with z/OS.
- Consider creating a new target work IODF based on the active IODF. During the inclusion of new devices failures might occur and the result would be an inconsistent target IODF. In such a case the user can easily start over by discarding the previously created target IODF and the active IODF always remains consistent.
- All active IODFs for the systems in a sysplex should be the same. This however, is not enforced. Consistent active IODFs allows IOS to avoid running some phases of discovery on every system. Tokens should be in sync and prior activates should be completed.
- Without indicating force full mode discovery, there is a limit on the number of successive failures. Unless force full mode discovery is requested, processing assumes that CUADD (logical address for a control unit) values start at 00 and continue through nn, with no missing CUADDs. Discovery for a controller ends after several successive failures occur and force full mode discovery is not indicated.
- For DASD, all newly discovered devices are assumed to be of type 3390 (either 3390B or 3390A types).
- If port restrictions (zoning or Prohibit Dynamic Connectivity Mask (PDCM)) exist within a switch (that is, they limit the ability of a CHPID to connect to a destination port for a control unit interface), I/O Autoconfiguration might configure paths that cannot be used. If a port is discovered on a controller, the assumption is that it has access to all configured logical control units on that controller.
- All CHPIDs, switches, and ports should be configured online and be accessible when discovery is attempted. With z/OS V2.1 you can exclude switches and CHPIDs from I/O Autoconfiguration.
- At least one system per CPC must be able to run dynamic I/O configuration changes and be part of the current sysplex. It must be able to make I/O configuration changes on behalf of the LP group systems on the CPC. This system does not have to be in the target LP group.
- A logical control unit that contains only secondary devices in an active peer-to-peer remote copy (PPRC) relationship might not be able to be discovered. The I/O used to determine the devices that are configured on a logical control unit cannot be run on secondary devices.
- I/O Autoconfiguration is a configuration tool that configures for availability. You can use Dynamic CHPID Management (DCM) for performance management. CHPID/path selection is run to minimize or even eliminate single points of failure to newly discovered logical control units. DCM manages for performance by adding CHPIDs and managing paths to the logical control units as needed.
- Within a target LPAR Group, I/O Autoconfiguration allows only controllers that are consistently defined (or absent) for the target LP group systems in the target IODF. If a controller is partially defined in the LP group (some systems have logical control units and devices configured that others do not have), I/O Autoconfiguration does not attempt to propose definitions for the systems within the LP group that do not have the definitions. You should control the discovery scope using LP groups with systems requiring the definition (that is, group systems in LP groups that share similar I/O requirements).

- If candidate access lists exclude an LPAR from accessing a control unit that is already defined on a CSS, I/O Autoconfiguration cannot discover and add that control unit. Therefore, all systems in the participating LP groups should have a homogeneous view of the devices and control units. If not, update the device candidate lists in HCD to add devices and control units to the LPARs.
- If switches are connected such that you can have three or more switches in a path to a control unit, this path can be chosen if no viable alternative exists.
- Run discovery attempts during times where changes are minimal. Do not make dynamic I/O configuration changes using HCD/HCM or the **`ACTIVATE`** IBM MVS™ command during the discovery process. ACTIVATE processing and CONFIGURE CHP commands will likely affect discovery processing.

For the most recent information, see the *z/OS HCD User's Guide* in the IBM Knowledge Center:

`http://www-01.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.cbdu100/iorules.htm%23iorules`

4.4 Overview of the I/O Autoconfiguration process

This section provides an overview of the I/O Autoconfiguration process.

The Fabric discovery process

The I/O Autoconfiguration function is started from the HCD Primary Task Selection panel. This action causes HCD to start IOS to run the Fabric discovery process. You can define the scope of discovery by searching all controllers, new controllers only, or search for the controller that contain a specific control unit.

Remember: The scope of the discovery is limited to the active sysplex.

The controller discovery

From the discovered controllers, HCD retrieves and proposes control unit and device types and numbers, channel path assignments, partition access, and OS device parameters. You can choose whether HCD runs the definition without user interaction, or whether the panel displays the proposed definitions so that you can confirm or change these values.

I/O Autoconfiguration makes temporary changes to the I/O configuration by adding devices that are used exclusively for discovery on the targeted systems to search for attached devices.

4.5 I/O Autoconfiguration using HCD

Running the I/O Autoconfiguration process using HCD involves these tasks, which are described in this section:

- Defining autoconfiguration policies
- Specifying profile options that are related to the zDAC process
- Defining logical partition groups for autoconfiguration
- Defining OS configuration groups for autoconfiguration
- Setting keywords for autoconfiguration policies
- Running the automatic I/O discovery process

- Applying updates to the autoconfiguration proposals
- Excluding a specific switch from the discovery and proposal process
- Excluding specific CHPIDs from the discovery and proposal process
- Before and after captures of the example

These instructions assume that you have a working knowledge of HCD.

4.5.1 Defining autoconfiguration policies

Before you use HCD to discover and define control units and I/O devices, specify your autoconfiguration policies. This task contains the following subtasks, which can be accessed from the HCD Primary Task menu, option 0, as shown in Figure 4-1:

- Option 0.1: HCD profile options
- Option 0.2: Autoconfiguration policies
- Option 0.3: LP groups for autoconfiguration
- Option 0.4: OS groups for autoconfiguration

```
                      z/OS V2.1 HCD
+--------------- Profile Options and Policies ----------------+
|                                                             |
|                                                             |
| Select type of data to define.                              |
|                                                             |
| __   1.  HCD profile options                                |
|      2.  Autoconfiguration policies                         |
|      3.  LP groups for autoconfiguration                    |
|      4.  OS groups for autoconfiguration                    |
|                                                             |
|                                                             |
|  F1=Help     F2=Split    F3=Exit     F9=Swap    F12=Cancel  |
+-------------------------------------------------------------+
```

Figure 4-1 HCD - Profile Options and Policies: options

Tip: If you want to restrict your discovery to a particular logical partition or group of logical partitions, define a Logical Partition Group first.

Similarly, with OS configurations define an OS Group first.

4.5.2 Specifying profile options that are related to the zDAC process

When I/O Autoconfiguration discovers and adds new device numbers to an OS configuration, HCD uses default OS config parameter values according to the device's unit information module (UIM).

These parameter default values can be overridden and set to your own preferences by selecting from the HCD Primary Option menu, option **0 - Edit profile options and policies**, then selecting option **1 - HCD profile options**.

Figure 4-2 on page 122 shows the parameter default values for device type 3390B.

```
+---------------- View Device Parameter / Feature Definition -----------------+
|                                                                 Row 1 of 6 |
| Command ===> ___________________________________________ Scroll ===> PAGE   |
|                                                                             |
| Configuration ID . : L06RMVS1       Sysplex systems                         |
| Device number  . . : 1000           Device type  . . . : 3390B              |
| Generic / VM device type  . . . . : 3390                                    |
|                                                                             |
| ENTER to continue.                                                          |
|                                                                             |
| Parameter/                                                                  |
| Feature    Value            R Description                                   |
| OFFLINE    Yes                Device considered online or offline at IPL    |
| DYNAMIC    Yes                Device supports dynamic configuration         |
| LOCANY     No                 UCB can reside in 31 bit storage              |
| WLMPAV     Yes                Device supports work load manager             |
| SHARED     Yes                Device shared with other systems              |
| SHAREDUP   No                 Shared when system physically partitioned     |
| ***************************** Bottom of data ****************************** |
|                                                                             |
|                                                                             |
|  F1=Help       F2=Split       F3=Exit       F7=Backward    F8=Forward       |
|  F9=Swap      F12=Cancel     F22=Command                                    |
+-----------------------------------------------------------------------------+
```

Figure 4-2 HCD - View Device Parameter: Default UIM values for device type 3390B

You might prefer to define your 3390B devices with `LOCANY=Yes` or `WLMPAV=No`, for example. To set these definitions, complete the following steps:

1. Select HCD option **0.1. HCD Profile Options**.
2. Page down until you find the profile option that says `OS_PARM_DEFAULT` (Figure 4-3).

```
Command ===> ___________________________________________ Scroll ===> CSR

Edit or revise profile option values.

HCD Profile : PBRESA.HCD.PROFILE

/ Profile keyword          A Value +
_ MAP_CUTYPE               Y ____________________________________________
# MCF_EXTENSION            Y 30
# MCF_VOL                  Y ____________________________________________
# MIGRATE_EXTENDED         Y NO
# MIXED_ESOTERIC           Y NO
_ OS_PARM_DEFAULT          Y ____________________________________________
# RCALL_LOG                Y NO
# RCALL_TIMEOUT            Y 60
# SHOW_IO_CHANGES          Y YES
# SHOW_CONFIG_ALL          Y NO
# SHOW_IOCP_DEFAULTS       N NO
```

Figure 4-3 HCD - HCD Profile Options: OS_PARM_DEFAULT not set

3. To enter an overriding value, type the OS parm you want to override, then press Enter (Figure 4-4).

```
 Command ===> ___________________________________________ Scroll ===> CSR

 Edit or revise profile option values.

 HCD Profile : PBRESA.HCD.PROFILE

 / Profile keyword          A Value +
 _ MAP_CUTYPE               Y ____________________________________________
 # MCF_EXTENSION            Y 30
 # MCF_VOL                  Y ____________________________________________
 # MIGRATE_EXTENDED         Y NO
 # MIXED_ESOTERIC           Y NO
 _ OS_PARM_DEFAULT          Y LOCANY,YES
 # RCALL_LOG                Y NO
 # RCALL_TIMEOUT            Y 60
 # SHOW_IO_CHANGES          Y YES
 # SHOW_CONFIG_ALL          Y NO
 # SHOW_IOCP_DEFAULTS       N NO
```

Figure 4-4 HCD - HCD Profile Options: OS_PARM_DEFAULT set

4. If you want to add more OS_PARM_DEFAULT changes, enter the letter **a** next to the existing `OS_PARM_DEFAULT` keyword entry and press Enter. An extra entry is added. Figure 4-5 shows entering `WLMPAV,NO`.

```
 Command ===> ___________________________________________ Scroll ===> CSR

Edit or revise profile option values.

HCD Profile : PBRESA.HCD.PROFILE

/ Profile keyword          A Value +
_ OS_PARM_DEFAULT          Y LOCANY,YES
_ OS_PARM_DEFAULT          Y WLMPAV,NO
# RCALL_LOG                Y NO
# RCALL_TIMEOUT            Y 60
# SHOW_IO_CHANGES          Y YES
# SHOW_CONFIG_ALL          Y NO
# SHOW_IOCP_DEFAULTS       N NO
# TSO_NOPREFIX             Y NO
# UNCOND_GENERATE_DROS     Y NO
# UIM_LIBNAME              N SYS1.NUCLEUS
# UIM_VOLSER               N ____________________________________________
```

Figure 4-5 HCD - HCD Profile Options: extra OS_PARM_DEFAULT set

5. If you now define new 3390B devices, the OS parameter defaults will now be overridden with what you specified in the HCD Profile Options (Figure 4-6).

```
---------------- View Device Parameter / Feature Definition ----------------
                                                                   Row 1 of 6
 Command ===> ___________________________________________ Scroll ===> PAGE

 Configuration ID . : L06RMVS1       Sysplex systems
 Device number  . . : 1000           Device type  . . . : 3390B
 Generic / VM device type  . . . . : 3390

 ENTER to continue.

 Parameter/
 Feature    Value            R Description
 OFFLINE    Yes                Device considered online or offline at IPL
 DYNAMIC    Yes                Device supports dynamic configuration
 LOCANY     Yes                UCB can reside in 31 bit storage
 WLMPAV     No                 Device supports work load manager
 SHARED     Yes                Device shared with other systems
 SHAREDUP   No                 Shared when system physically partitioned
 **************************** Bottom of data *****************************
```

Figure 4-6 HCD - View Device Parameter: default UIM values for device type 3390B overridden

Note: Default values *cannot* be set differently for different device types. Additionally, the assignment of devices to esoteric names cannot currently be controlled with an I/O Autoconfiguration (zDAC) dialog.

4.5.3 Defining logical partition groups for autoconfiguration

A logical partition group (LP group) is a collection of logical partitions in the same sysplex used by I/O Autoconfiguration to determine which partitions the discovered devices are assigned to.

To define LP groups for autoconfiguration, complete the following steps:

1. Select option **3. LP groups for autoconfiguration** from the Profile Options and Policies menu. This action starts the Autoconfiguration LP Group List, which displays a list of LP groups. In this example, no lists are defined.
2. Enter **add** in the command line and press Enter (Figure 4-7).

```
------------- Autoconfiguration LP Group List -------------

Command ===> add________________________ Scroll ===> CSR

To view assigned partitions, select one or more LP groups,
then press Enter. To add an LP group, use F11.

/ LP group name   Description
******************** Bottom of data *********************
```

Figure 4-7 HCD - Autoconfiguration LP Group List: adding an LP group

Use the Autoconfiguration LP Group List panel to define or delete LP groups, to assign logical partitions to a group, or unassign partitions from a group.

3. Enter your LP group name and description, and press Enter (Figure 4-8).

```
------------ Add Autoconfiguration LP Group ------------


Specify name and description for new LP group.

LP group name  . . . . . P501LPGP
Description  . . . . . . Test LP Group for SCZP501______
```

Figure 4-8 HCD - Add Autoconfiguration LP Group panel: defining new group name

4. Enter **s** next to the LP group name to add processors and partitions, and press Enter (Figure 4-9).

```
------------- Autoconfiguration LP Group List -------------
                                                 Row 1 of 1
Command ===> ___________________________ Scroll ===> CSR

To view assigned partitions, select one or more LP groups,
then press Enter. To add an LP group, use F11.

/ LP group name   Description
s P501LPGP        Test LP Group for SCZP501
********************* Bottom of data *********************
```

Figure 4-9 HCD - Autoconfiguration LP Group List: adding partitions

5. Enter **add** in the command line and press Enter or press PF11 (Figure 4-10).

```
--------- Autoconfiguration LP Group Assignment List ----------

Command ===> add____________________________ Scroll ===> CSR

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

LP group name : P501LPGP     Test LP Group for SCZP501

/ Partition Name       Description
*********************** Bottom of data ************************
```

Figure 4-10 HCD - Autoconfiguration LP Group Assignment List: adding processors and partitions

6. Enter the Processor ID and Partition Name and press Enter. For a list of Processor IDs and Partitions that are defined in the IODF, press PF4 in the data fields (Figure 4-11).

```
---------------------- Add Partition to LP Group -------------------


Specify the following values.

LP group name  . . . . . : P501LPGP     Test LP Group for SCZP501

Processor ID . . . . . . . SCZP501   +
Partition Name . . . . . . A01       +
```

Figure 4-11 HCD - Add Partition to LP Group: specifying a Processor ID and Partition Name

Processor SCZP501 and LPAR A01 of CSS.0 are added to the LP Group P501LPGP (Figure 4-12).

```
--------- Autoconfiguration LP Group Assignment List ----------
                                                     Row 1 of 1
Command ===> _______________________________ Scroll ===> CSR

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

LP group name : P501LPGP     Test LP Group for SCZP501

/ Partition Name       Description
_ SCZP501.0.A01        SC76
********************** Bottom of data ***********************
```

Figure 4-12 HCD - Autoconfiguration LP Group Assignment List: Processor and partition added

7. More processors and partitions for this sysplex can be added by entering **add** in the command line or by entering the letter **a** next to an existing entry (Figure 4-13).

```
--------- Autoconfiguration LP Group Assignment List -----------
                                                     Row 1 of 2
Command ===> _______________________________ Scroll ===> CSR

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

LP group name : P501LPGP     Test LP Group for SCZP501

/ Partition Name       Description
_ SCZP401.0.A08        WTSCPLX8 SC80
_ SCZP501.0.A01        SC76
*********************** Bottom of data ************************
```

Figure 4-13 HCD - Autoconfiguration LP Group Assignment List: adding more processors and partitions

You can now return to the Profile Options and Policies menu.

4.5.4 Defining OS configuration groups for autoconfiguration

An OS group is a collection of OS configurations. It is used by I/O Autoconfiguration to determine to which operating systems of type "MVS" the auto-defined devices are assigned.

To define OS configuration groups for autoconfiguration, complete the following steps:

1. Select option **4. OS groups for autoconfiguration** from the Profile Options and Policies menu (Figure 4-14). This action starts the Autoconfiguration OS Group List, which displays a list of operating system groups (OS groups).
2. Type **add** in the command line and press Enter or press PF11.

```
------------- Autoconfiguration OS Group List -------------

Command ===> add________________________ Scroll ===> CSR

To view assigned OS configurations, select one or more OS
groups, then press Enter. To add an OS group, use F11.

/ OS group   Description
********************* Bottom of data *********************
```

Figure 4-14 HCD - Autoconfiguration OS Group List: adding a OS group

Use the Autoconfiguration OS Group List to view, add, or delete OS groups, or to assign autoconfigured devices to operating systems.

3. Type your group name and description, and press Enter (Figure 4-15).

```
------------ Add Autoconfiguration OS Group ----------


Specify name and description for new OS group.

OS group name  . . . . . P501OSGP
Description  . . . . . . Test OS Group for SCZP501___
```

Figure 4-15 HCD - Add Autoconfiguration OS group: OS group name added

4. Enter **s** next to the OS group name to add one or more OS configuration groups (Figure 4-16).

```
------------ Autoconfiguration OS Group List -------------
                                                 Row 1 of 1
Command ===> ___________________________ Scroll ===> CSR

To view assigned OS configurations, select one or more OS
groups, then press Enter. To add an OS group, use F11.

/ OS group   Description
s P501OSGP   Test OS Group for SCZP501
******************** Bottom of data *********************
```

Figure 4-16 HCD - Add Autoconfiguration OS group: select an OS group to add an OS Config to

5. Enter **add** in the command line and press Enter (Figure 4-17).

```
--------- Autoconfiguration OS Group Assignment List ----------

Command ===> add____________________________ Scroll ===> CSR

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

OS group name : P501OSGP     Test OS Group for SCZP501

/ OS Configuration ID  Description
********************** Bottom of data ***********************
```

Figure 4-17 HCD - Add Autoconfiguration OS group: adding an OS Config to an OS group

6. Enter your OS config ID, or press PF4 in the OS configuration ID data field to be prompted for the OS configurations that are defined in the IODF and press Enter (Figure 4-18).

```
----------- Add Operating System Configuration to OS Group ---------


Specify the following values.

OS group name . . . . : P501OSGP     Test OS Group for SCZP501

OS configuration ID . . L06RMVS1  +
```

Figure 4-18 HCD - Add Operating system Configuration to OS Group: selecting the OS Config

The OS configuration ID L06RMVS1 has been added to the OS group (Figure 4-19).

```
--------- Autoconfiguration OS Group Assignment List ----------
                                                     Row 1 of 1
Command ===> _______________________________ Scroll ===> CSR

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

OS group name : P501OSGP     Test OS Group for SCZP501

/ OS Configuration ID  Description
_ L06RMVS1             Sysplex systems
*********************** Bottom of data ************************
```

Figure 4-19 HCD - Autoconfiguration OS Group Assignment List: OS Config added to OS group

7. More OS configuration IDs for this sysplex can be added by entering `add` on the command line or entering `a` next to the existing entry.

You can now return to the Profile Options and Policies menu.

4.5.5 Setting keywords for autoconfiguration policies

To set the autoconfiguration policies by using keywords, select option **2. Autoconfiguration policies** from the Profile Options and Policies menu. This action opens the Autoconfiguration Policies panel (Figure 4-20 on page 130).

This panel lists the current value settings for autoconfiguration policy keywords as they are either explicitly set in the HCD profile data set or defaulted by HCD. Use this panel to revise or change the displayed keyword values.

```
----------------------- Autoconfiguration Policies ------------------------
                                                  Row 1 of 13 More:       >
Command ===> ___________________________________________ Scroll ===> CSR

Edit or revise autoconfiguration policies.

HCD Profile : PBRESA.HCD.PROFILE

/ Policy keyword           P Value +
_ AUTO_CHPID_EXCLUDE       N ____________________________________________
_ AUTO_CHPID_INCLUDE       N ____________________________________________
# AUTO_MATCH_CU_DEVNUM     Y YES
# AUTO_SS_ALTERNATE        Y 1
# AUTO_SS_DEVNUM_SCHEME    Y PAIRING
# AUTO_SUG_CU_RANGE        Y 0001-FFFE
# AUTO_SUG_DEV_RANGE       Y 0001-FFFF
# AUTO_SUG_DYN_CHPIDS      Y 2
# AUTO_SUG_LPGROUP         N ____________________________________________
# AUTO_SUG_OSGROUP         N ____________________________________________
# AUTO_SUG_STAT_CHPIDS     Y 6
_ AUTO_SWAD_EXCLUDE        N ____________________________________________
_ AUTO_SWAD_INCLUDE        N ____________________________________________
***************************** Bottom of data ******************************
```

Figure 4-20 HCD - Autoconfiguration policies: default values

You must set the following policy keywords:

- AUTO_MATCH_CU_DEVNUM: This policy specifies whether, for autoconfiguration definitions, a control unit number must match the starting base device number.

  If you specify **YES** (which is the *default*), the first base device is set to the same number as the control unit. If NO is specified, the device number of the first base device and the control unit number do not necessarily need to match.

- AUTO_SS_ALTERNATE: This policy specifies the ID of the subchannel set in which newly discovered parallel access volume (PAV) alias devices are defined during an auto-definition process. Free device numbers must be available in this subchannel set, and processors that have access to the device range must support alternate subchannels.

  The *default* subchannel set ID is 1.

- AUTO_SS_DEVNUM_SCHEME: This policy defines the scheme for assigning device numbers to PAV alias devices in an alternate subchannel set. The following schemes are supported:
  - CONSECUTIVE: The alias device numbers in an alternate subchannel set are consecutive to the base device numbers.
  - DENSE: The device numbers in an alternate subchannel set are densely assigned, that is, the next free device numbers in the assigned device number range are used.
  - PAIRING: This is the *default*. Base and alias device numbers are assigned alternatively starting with, for example, device numbers xx00 and xx80 versus xx80 and xx00.
  - NONE: The control unit and device number proposals by HMC are bypassed. This lets the user manually apply the numbers for detected objects. (new with z/OS V2R1).

- AUTO_SUG_CU_RANGE: This policy specifies the range of control unit numbers from which numbers for auto-defined control units are taken. If no value is specified, the range 0001 - FFFE is taken as *default*. Specify the range according to the following syntax:

  nnnn-mmmm

  Where:

  – nnnn is the lower range boundary.
  – mmmm is the upper range boundary.

- AUTO_SUG_DEV_RANGE: This policy specifies the range of device numbers from which device numbers for auto-defined devices are taken. If no value is specified, the range 0001 - FFFF is taken as *default*.

  Consideration: I/O Autoconfiguration avoids using device and control unit numbers in the 0000-00FF range in Subchannel set 0. If you have no alternative number ranges available, you must configure the CUs or devices manually in this range.

- AUTO_SUG_DYN_CHPIDS: This policy specifies the number of dynamically managed channel paths that are allowed on a control unit definition, if it is auto-defined. A maximum number of seven dynamic channel paths is allowed. The *default* is 2.
- AUTO_SUG_LPGROUP: This policy specifies the name of a group of logical partitions to which discovered devices are assigned. If no name is set, devices are assigned to all eligible partitions of the active sysplex. The reserved group name ALL signals this setting during the autoconfiguration processing.
- AUTO_SUG_OSGROUP: This policy specifies the name of a group of OS configurations to which discovered devices are assigned. If no name is set, devices are assigned to all OS configurations that correspond to the active LP group.
- AUTO_SUG_STAT_CHPIDS: This policy specifies the number of static channel paths to be assigned to a control unit definition, if it is auto-defined. At least one and not more than eight static channel paths can be defined. The *default* is 6.

The next four keywords became available with z/OS V2R1.

- AUTO_CHPID_EXCLUDE: This policy specifies the exclusion of CHPIDs in the discovery and proposal process. To add a CHPID for exclusion enter the letter **a** next to the keyword to add a CHPID, or remove a previously added CHPID with the letter **d**. This applies also for all four keywords.
- AUTO_CHPID_INCLUDE: This policy specifies the inclusion of CHPIDs into the discovery and proposal process.
- AUTO_SWAD_EXCLUDE: This policy specifies the exclusion of switches in the discovery and proposal process.
- AUTO_SWAD_INCLUDE: This policy specifies the inclusion of switches into the discovery and proposal process.

4.5.6 Running the automatic I/O discovery process

After you specify the policies, start the I/O Autoconfiguration process for HCD to discover and automatically define control units and I/O devices into a specified target IODF.

Start the I/O Autoconfiguration process from HCD by completing the following steps:

1. Select option **1. Define, modify, or view configuration data** from the Primary Task Selection menu.
2. Select option **6. Discovered new and changed control units and I/O devices**.

These steps open the Discovery and Autoconfiguration Options panel (Figure 4-21). Use this panel to select processing options for discovery and autoconfiguration. Input to all fields is required, and the initial defaults are supplied.

```
+----------------- Discovery and Autoconfiguration Options -----------------+
|                                                                           |
|                                                                           |
| Specify autoconfiguration options. Then, press Enter to start the         |
| discovery process.                                                        |
|                                                                           |
| Autoconfiguration is based on 2  1. Active IODF                           |
|                                  2. Currently accessed IODF               |
|                                                                           |
| Scope of discovery  . . . . . 1  1. New controllers only                  |
|                                  2. All controllers                       |
|                                  3. Controller containing CU ____    +    |
|                                  4. Controller with S/N      ____________ |
|                                                                           |
| Show proposed definitions . . 1  1. Yes   2. No                           |
| Force full mode discovery . . 2  1. Yes   2. No                           |
| Tolerate incapable systems    2  1. Yes   2. No                           |
|                                                                           |
| Target IODF name  . . . 'SYS6.IODFC5.WORK'                         +      |
|                                                                           |
```

Figure 4-21 HCD - Discovery and Autoconfiguration Options: setting discovery options

You must set the following autoconfiguration options:

- Autoconfiguration is based on

  Choose whether the active production IODF or the currently accessed IODF is used as the base for new configuration definitions that result from the discovery process.

  HCD checks the discovered devices accessible to the system against this selected IODF, whether they are already defined or not. If they are found to be new or changed, the resulting configuration proposals are also adopted to fit into this IODF.

  HCD copies this IODF to the selected target IODF, which receives all changes that are made to the configuration during autoconfiguration processing.

  > **Remember:** Any IODF that you select as the base for I/O autoconfiguration must be enabled for full dynamic activation. If you select a work IODF that is not enabled, base it on a production IODF.

- Scope of discovery

  In this selection, you decide which controllers are discovered:

  – New controllers only: HCD discovers and returns only new controllers that are not yet known in the chosen target IODF.

  – All controllers: HCD discovers and returns all new controllers and all changed controllers (a full discovery).

  – Controller containing CU ____: HCD runs a discovery that is limited to the controller containing the control unit with the specified number. The referenced control unit must be a DASD or tape CU, and must be defined in the target IODF.

  – Controller with S/N: HCD runs a discovery that is limited to the controller containing the serial number with the specified number.

- Show proposed definitions

  You can decide whether the panel displays proposed definitions for possible configuration changes. Select `Yes` if you want to work in an attended operation mode. In this mode, HCD starts a subsequent series of panels in which you can revise and change the proposed settings.

  Select **No** if you want to run the unattended fast path of I/O Autoconfiguration. In this case, HCD does not offer a possibility to revise the proposals, or to update or add definitions. Instead, the HCD definitions are saved in the target IODF immediately.

- Force full mode discovery

  Decide when discovery processing should stop. If set to `No`, which is the default, processing stops after several consecutive unused CUADD values that do not exist on the target controller. With this option set to `Yes`, for each discovered controller, all unused logical control unit addresses (CUADD values) and unit addresses are checked for changes.

- Tolerate incapable systems (This definition became available with z/OS V2R1)

  Select `Yes` to decide to tolerate incapable systems. I/O Autoconfiguration allows discovery by controller serial number and filters the discovered controllers accordingly. HCD can process an I/O Autoconfiguration request that is partially directed against unavailable systems of an LPAR group or a sysplex, or against systems that are not capable to support I/O Autoconfiguration. Users can specify that the request applies to appropriate systems only, and that unavailable/incapable systems are tolerated but ignored.

- Target IODF name

  Type the name of a work IODF to receive the configuration definitions for all discovered new or changed controllers, according to your selected scope of discovery.

  This input is required. The specified IODF can either be an existing work IODF, or created by HCD. In any case, the IODF specified in the `Autoconfiguration is based on` field is copied to the specified target IODF.

  The target IODF must not be enabled for multi-user access.

  > **Tip:** As soon as you accepted any proposals into your target IODF, it becomes the new currently accessed IODF.

3. Press Enter and observe the status message at the bottom of the panel (Figure 4-22).

```
+----------------------------------------------+
| FABRIC discovery in progress - please wait ... |
+----------------------------------------------+
```

Figure 4-22 HCD - Fabric Discovery in progress message

4.5.7 Applying updates to the autoconfiguration proposals

After a successful discovery, HCD displays the result in the Discovered New or Changed Controller List (Figure 4-23). Only discovered controllers are returned, which are reachable from all target systems that have partitions that are defined in the LP group referenced by the AUTO_SUG_LPGROUP policy keyword.

```
+--------- Discovered New or Changed Controller List ---------+
|   Policy  Backup  Query  Help                                |
| ----------------------------------------------------------  |
|                                                  Row 1 of 3 |
| Command ===> ___________________________ Scroll ===> CSR    |
|                                                              |
| Select one or more controllers to be defined, then press    |
| Enter.                                                       |
|                                                              |
|                  Manufacturer                                |
| /  Type   Model  Name  Plant   Serial-#     New  Processed  |
| _  2107   932    IBM   75      FBAZ1        Yes  No         |
| _  3590   J70    IBM   78      C4146        Yes  No         |
| _  3590   C06    IBM   78      C5367        Yes  No         |
| ******************** Bottom of data ********************** |
```

Figure 4-23 HCD - Discovered New or Changed Controller List: discovered controllers

This panel lists all discovered controllers that are either not yet defined in the IODF or whose definition in the IODF differs from discovered controller characteristics.

To apply updates to the autoconfiguration proposals, complete the following steps:

1. Enter the action code “/” next to one or multiple controllers that you want to be defined or changed in the target IODF. HCD then processes each selected controller in the way that is described below.
2. Press Enter and note the status message at the bottom of the panel (Figure 4-24).

```
+--------------------------------------------------+
| CONTROLLER discovery in progress - please wait ... |
+--------------------------------------------------+
```

Figure 4-24 HCD - Controller discovery in progress

3. In this example, HCD issued msg CBDG721I, indicating that the serial number field for the following control units has been updated in the IODF with the serial number that was discovered (Figure 4-25). Press PF3 to continue.

```
----------------------- Proposed Control Unit List ------------------------

------------------------------ Message List -------------------------------
  Save  Query  Help
--------------------------------------------------------------------------
                                                                Row 1 of 4
Command ===> ___________________________________________ Scroll ===> CSR

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

/ Sev Msg. ID  Message Text
_ I   CBDG721I Discovery provided serial number FBAZ1 for controller.
#              The serial numbers of the following control units are
#              changed accordingly: 6100, 6300, 6500, 6700, 6000, 6200,
#              6400, 6600
**************************** Bottom of data *****************************
```

Figure 4-25 HCD - Proposed Control Unit List: CBDG721I message for serial number changes

4. As a result of the discovery process, the Proposed Control Unit List (Figure 4-26) offers definition proposals for the control units that are found in the currently processed controller.

```
----------------------- Proposed Control Unit List ------------------------
                                                                Row 1 of 8
Command ===> ___________________________________________ Scroll ===> CSR

Control unit type  . . : 2107-932       Serial number  : FBAZ1

Proposed switch.ports  : 62.85 62.58 61.08 62.0C 61.0C 61.58 62.08 61.85

To accept the proposed values, press Enter. To modify them, edit the
fields, or select one or more control units to change, exclude or include
the corresponding definitions, then press Enter.


  CU   CU       # of     LPAR
/ ADD  number+  devices  Access+   New Description                      I
_ 00   6000     42       ALL       No  YY CU 0                          Y
_ 01   6100     42       ALL       No  YY CU 1                          Y
_ 02   6200     42       ALL       No  YY CU 2                          Y
_ 03   6300     42       ALL       No  YY CU 3                          Y
_ 04   6400     42       ALL       No  YY CU 4                          Y
_ 05   6500     42       ALL       No  YY CU 5                          Y
_ 06   6600     42       ALL       No  YY CU 6                          Y
_ 07   6700     42       ALL       No  YY CU 7                          Y
**************************** Bottom of data *****************************
```

Figure 4-26 HCD - Proposed Control Unit List: accept, modify or not

You can accept the proposed control unit definitions, or make the following modifications:

– For control units that show `Yes` in column New (which indicates whether the control unit is not yet defined in the IODF), you can overwrite the values in column CU number.

– Also, you can overwrite the LPAR Access and Description fields by entering values in the panel.

– In the forward slash (/) column, you can enter an action code (i, e, or c) for the given control unit:

  - With action code **i**, you can include the corresponding control unit definition in the IODF.
  - With action code **e**, you can exclude the control unit from being defined in the IODF.
  - Your selection is reflected in column **I.** Y denotes included and N denotes excluded control units.
  - Action code **c** leads you to the Select Processor / CU panel. On this panel, HCD displays a list of all defined processors. You can define how the control unit is to be attached to one or more processors.

Rules for discovered control units:

– For each discovered control unit that is already defined with the same CUADD value, the existing control unit definition is checked for the same serial number. If the serial numbers match, or the IODF definition does not contain a serial number, the control unit number of the existing control unit is used. If the serial numbers do not match, a warning message is given, and the discovered control unit is proposed with a new number.

– For each discovered control unit that is not yet defined in the IODF, a new serial number is proposed.

– When a new control unit number is proposed, its value is taken from the preferred range that is specified by the AUTO_SUG_CU_RANGE policy. If there is no free control unit number in the IODF within that range, a warning message indicates that the policy cannot be followed. A free control unit number outside of the range is then proposed.

– Proposed existing control units are updated with the discovered serial number. If the type of a discovered control unit differs from its definition in the IODF, the definition is updated.

Rules for discovered devices:

– For each discovered device that is already defined with the same unit address on an existing control unit, the existing device number is proposed. For non-existing devices on the control unit, the existing device numbering scheme is applied if possible.

– For new devices on new or existing control units where the existing device number scheme cannot be applied, the device numbers are determined based on the AUTO_SUG_DEV_RANGE and AUTO_MATCH_CU_DEVNUM policies.
For PAV alias devices, the numbers are additionally determined based on the AUTO_SS_ALTERNATE and AUTO_SS_DEVNUM_SCHEME policies. If a policy cannot be applied because no free numbers are available for the active LP group and OS groups, a warning message is given. Free device numbers outside the policies can then be used.

5. Apply the modifications and press Enter when you finish, or you can accept the proposed definitions without changes and press Enter. In both cases, HCD displays the Proposed Control Unit / Device List (Figure 4-27).

```
                     Proposed Control Unit / Device List          Row 1 of 23
Command ===> _______________________________________________ Scroll ===> CSR

Control unit type  . . : 2107-932       Serial number  : FBAZ1

To accept the proposed values, press Enter. To modify them, edit the
fields, or select one or more device ranges to change, exclude or include
the corresponding definitions, then press Enter.


  ------Device----- S CU   UA    OS
/ Number   Type+    S Num  Range Access+  N Description                      I
_ 5F3A,10  3390B    0 6000 00-09 ALL      N ________________________________ Y
_ 5F45,8   3390B    0 6000 0A-11 ALL      N ________________________________ Y
_ 5F55,2   3390B    0 6000 12-13 ALL      N ________________________________ Y
_ 5F58,12  3390B    0 6000 14-1F ALL      N ________________________________ Y
_ 5F65,5   3390B    0 6000 20-24 ALL      N ________________________________ Y
_ 5F70     3390B    0 6000 25-25 ALL      N ________________________________ Y
_ 5F72,4   3390B    0 6000 26-29 ALL      N ________________________________ Y
_ 5F78,12  3390B    0 6100 00-0B ALL      N ________________________________ Y
_ 5F85     3390B    0 6100 0C-0C ALL      N ________________________________ Y
_ 5F8F,4   3390B    0 6100 0D-10 ALL      N ________________________________ Y
_ 5F94     3390B    0 6100 11-11 ALL      N ________________________________ Y
_ 5FA7,2   3390B    0 6100 23-24 ALL      N ________________________________ Y
_ 5FAA,5   3390B    0 6100 25-29 ALL      N ________________________________ Y
_ 5FB8,17  3390B    0 6100 12-22 ALL      N ________________________________ Y
_ 5FAF,9   3390B    0 6200 00-08 ALL      N ________________________________ Y
_ 5FC9,3   3390B    0 6200 11-13 ALL      N ________________________________ Y
_ 5FCC,8   3390B    0 6200 09-10 ALL      N ________________________________ Y
_ 6214,22  3390B    0 6200 14-29 ALL      N ________________________________ Y
_ 6300,42  3390B    0 6300 00-29 ALL      N ________________________________ Y
_ 6400,42  3390B    0 6400 00-29 ALL      N ________________________________ Y
_ 6500,42  3390B    0 6500 00-29 ALL      N ________________________________ Y
_ 6600,42  3390B    0 6600 00-29 ALL      N ________________________________ Y
_ 6700,42  3390B    0 6700 00-29 ALL      N ________________________________ Y
****************************** Bottom of data ********************************
```

Figure 4-27 HCD - Proposed Control Unit / Device List: proposed values

This list proposes definition details for existing or new devices accessible by the currently processed discovered control units (in the example, 6000 - 6700). In the header of this panel, you can see the control unit type and serial number of the discovered controller.

You can accept the proposed device definitions without changes by pressing Enter. Also, you can narrow the proposed device definitions by overwriting one or more of the device ranges. You can only overwrite devices with a **Y** in column N (abbreviation for New), which indicates that the device range is not yet defined in the IODF.

Furthermore, for one or more of the listed device ranges with **Y** in column N, you can change the OS Access and the Description fields by overwriting the values in the panel. Again, a changed OS group must be a subset of the initial OS group.

For further available actions on devices, select one or more devices using the forward slash (/) action code:

– Selecting action code `i` includes the corresponding devices from autoconfiguration, Similarly, action code `e` excludes them.

– Using action code `c` leads you to the Device / Processor Definition panel. This panel displays a list of all defined processors that have one or more channel paths to the control unit to which the device being added or changed is attached. Here you can select the processor and channel subsystem for which you want to change the device-to-processor definition.

In both cases, either with modifications applied or with accepting the unchanged propositions, pressing Enter returns you to the Discovered New or Changed Controller List.

6. In the example, accept the discovered devices and press Enter.

   For each successfully processed controller, its Processed field is now turned to **Yes** (Figure 4-28).

```
-------- Discovered New or Changed Controller List --------
  Policy  Backup  Query  Help
----------------------------------------------------------
                                                Row 1 of 3
Command ===> ___________________________ Scroll ===> CSR

Select one or more controllers to be defined, then press
Enter.

                 Manufacturer
/  Type   Model  Name  Plant   Serial-#     New  Processed
_  2107   932    IBM   75      FBAZ1        Yes  Yes
_  3590   J70    IBM   78      C4146        Yes  No
_  3590   C06    IBM   78      C5367        Yes  No
********************* Bottom of data *********************
```

Figure 4-28 HCD - Discovered New or Changed Controller List: discovered controller added

You can select the next controller to be autoconfigured, or press PF3 to exit the panel.

4.5.8 Excluding a specific switch from the discovery and proposal process

In this section, we describe how to work with the keyword **AUTO_SWAD_EXCLUDE** in the Autoconfiguration Policies panel.

Figure 4-29 is an excerpt showing proposed switch ports for switch 61 and switch 62 of the previous controller discovery on control unit 2107-932 with serial 75-FBAZ1.

```
----------------------- Proposed Control Unit List ------------------------
                                                                  Row 1 of 8
Command ===> ___________________________________________ Scroll ===> CSR

Control unit type  . . : 2107-932       Serial number  : FBAZ1

Proposed switch.ports  : 62.85 62.58 61.08 62.0C 61.0C 61.58 62.08 61.85
```

Figure 4-29 HCD - Proposed Control Unit List: switch ports for switch 61 and switch 62

Figure 4-30 shows the Autoconfiguration Policies panel with the AUTO_SWAD_EXCLUDE keyword and a value of 61 to exclude all proposed switch ports related to switch 61.

```
------------------------ Autoconfiguration Policies ------------------------
                                                    Row 3 of 13 More:       >
 Command ===> ___________________________________________ Scroll ===> CSR

 Edit or revise autoconfiguration policies.

 HCD Profile : PBRESA.HCD.PROFILE

 / Policy keyword          P Value +
 # AUTO_MATCH_CU_DEVNUM    Y YES
 # AUTO_SS_ALTERNATE       Y 1
 # AUTO_SS_DEVNUM_SCHEME   Y PAIRING
 # AUTO_SUG_CU_RANGE       Y 0001-FFFE
 # AUTO_SUG_DEV_RANGE      Y 0001-FFFF
 # AUTO_SUG_DYN_CHPIDS     Y 2
 # AUTO_SUG_LPGROUP        N ____________________________________________
 # AUTO_SUG_OSGROUP        N ____________________________________________
 # AUTO_SUG_STAT_CHPIDS    Y 6
 _ AUTO_SWAD_EXCLUDE       N 61
 _ AUTO_SWAD_INCLUDE       N ____________________________________________
```

Figure 4-30 HCD - Autoconfiguration Policies: AUTO_SWAD_EXCLUDE set for switch 61

You can add multiple AUTO_SWAD_EXCLUDE entries for more than just one switch to be excluded. To enter the next value type **a** before the keyword press enter and add the next value to the new line. See Figure 4-31

```
---------------------- Proposed Control Unit List -----------------------
                                                                Row 1 of 8
Command ===> ___________________________________________ Scroll ===> CSR

Control unit type  . . : 2107-932      Serial number  : FBAZ1

Proposed switch.ports  : 62.85 62.58 62.0C 62.08

To accept the proposed values, press Enter. To modify them, edit the
fields, or select one or more control units to change, exclude or include
the corresponding definitions, then press Enter.


  CU   CU       # of     LPAR
/ ADD  number+  devices  Access+   New Description                       I
_ 00   6000     42       ALL       No  YY CU 0                           Y
_ 01   6100     42       ALL       No  YY CU 1                           Y
_ 02   6200     42       ALL       No  YY CU 2                           Y
_ 03   6300     42       ALL       No  YY CU 3                           Y
_ 04   6400     42       ALL       No  YY CU 4                           Y
_ 05   6500     42       ALL       No  YY CU 5                           Y
_ 06   6600     42       ALL       No  YY CU 6                           Y
_ 07   6700     42       ALL       No  YY CU 7                           Y
***************************** Bottom of data ******************************
```

Figure 4-31 HCD - Proposed Control Unit List: without switch port related to switch 61

This Autodiscovery is based on the same starting conditions (work IODF) as the two discovery sequences shown in this section.

4.5.9 Excluding specific CHPIDs from the discovery and proposal process

In this section, we describe how use the `AUTO_CHPID_EXCLUDE` keyword in the Autoconfiguration Policies panel.

Figure 4-32 on page 141 shows the Autoconfiguration Policies panel with a new setup. The exclusion of switch 61 was removed. Instead, all CHPIDs attached to switch 62 are excluded now.

The syntax for the value is Processor.CSS,CHPID. In our example `SCZP501.0,41` for the first CHPID 41 in channel subsystem 0. In addition, a wildcard with asterisk (*) and a range are allowed also.

```
------------------------ Autoconfiguration Policies ------------------------
                                                  Row 1 of 21 More:       >
 Command ===> ___________________________________________ Scroll ===> CSR

 Edit or revise autoconfiguration policies.

 HCD Profile : PBRESA.HCD.PROFILE

 / Policy keyword            P Value +
 _ AUTO_CHPID_EXCLUDE        N SCZP501.0,41
 _ AUTO_CHPID_EXCLUDE        N *.*,43
 _ AUTO_CHPID_EXCLUDE        N *.*,45
 _ AUTO_CHPID_EXCLUDE        N *.*,47
 _ AUTO_CHPID_EXCLUDE        N *.*,4E-4F
 _ AUTO_CHPID_EXCLUDE        N *.*,54-57
 _ AUTO_CHPID_EXCLUDE        N *.*,5C-5F
 _ AUTO_CHPID_EXCLUDE        N *.*,63
 _ AUTO_CHPID_EXCLUDE        N *.*,65
 _ AUTO_CHPID_INCLUDE        N ____________________________________________
 # AUTO_MATCH_CU_DEVNUM      Y YES
 # AUTO_SS_ALTERNATE         Y 1
 # AUTO_SS_DEVNUM_SCHEME     Y PAIRING
 # AUTO_SUG_CU_RANGE         Y 0001-FFFE
 # AUTO_SUG_DEV_RANGE        Y 0001-FFFF
 # AUTO_SUG_DYN_CHPIDS       Y 2
 # AUTO_SUG_LPGROUP          N ____________________________________________
 # AUTO_SUG_OSGROUP          N ____________________________________________
 # AUTO_SUG_STAT_CHPIDS      Y 6
 _ AUTO_SWAD_EXCLUDE         N ____________________________________________
 _ AUTO_SWAD_INCLUDE         N ____________________________________________
***************************** Bottom of data ******************************
```

Figure 4-32 HCD - Autoconfiguration Policies: AUTO_CHPID_EXCLUDE for all CHPIDs attached to switch 62

4.5.10 Before and after captures of the example

The following panels of the control unit list (Figure 4-33 on page 142) and control unit definition (Figure 4-34 on page 142) are used in the example before the I/O Autoconfiguration process. They show the control unit definitions that are currently connected to processors other than the processor in the LP Group.

Note the following items:

- The CSS column indicates 6.
- The MC column is blank.
- The Serial # is blank.

Figure 4-33 shows the control unit list before the changes.

```
                              Control Unit List                Row 718 of 904
Command ===> ______________________________________________ Scroll ===> CSR

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

                           ---#---
/ CU   Type +        CUADD CSS MC  Serial-# + Description
_ 6000 2107          0     6       __________ YY CU 0
_ 6100 2107          1     6       __________ YY CU 1
_ 6200 2107          2     6       __________ YY CU 2
_ 6300 2107          3     6       __________ YY CU 3
_ 6400 2107          4     6       __________ YY CU 4
_ 6500 2107          5     6       __________ YY CU 5
_ 6600 2107          6     6       __________ YY CU 6
_ 6700 2107          7     6       __________ YY CU 7
```

Figure 4-33 HCD - Control Unit List: example before I/O Autoconfiguration

Figure 4-34 shows the definition before the changes.

```
----------------------- View Control Unit Definition -----------------------
                                                    Row 1 of 6 More:      >
Command ===> ____________________________________________ Scroll ===> CSR

Control unit number . : 6000          YY CU 0
Control unit type . . : 2107          Serial number . . . :

Connected switch.ports: 61.08 61.58 62.08 62.58


ENTER to continue.

           ----------------Channel Path ID . Link Address-----------------
Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
SCZP201.0  40.08   42.58   41.08   43.58
SCZP201.1  40.08   42.58   41.08   43.58
SCZP401.0  61.08   64.58   62.08   65.58
SCZP401.1  61.08   64.58   62.08   65.58
SCZP401.2  61.08   64.58   62.08   65.58
SCZP401.3  61.08   64.58   62.08   65.58
****************************** Bottom of data ******************************
```

Figure 4-34 HCD - View Control Unit Definition: example before I/O Autoconfiguration

Figure 4-35 on page 143 shows the panel of the control unit list that is used in the example after the I/O Autoconfiguration process that shows the CU definitions.

Note the following changes:

- The CSS column now indicates 7.
- The MC column now indicates 2.
- The Serial # now indicates the Serial Number of the controller.

```
                              Control Unit List                 Row 718 of 904
Command ===> ______________________________________________ Scroll ===> CSR

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

                           ---#---
/ CU   Type +        CUADD CSS MC  Serial-# + Description
_ 6000 2107          0     7   2   FBAZ1      YY CU 0
_ 6100 2107          1     7   2   FBAZ1      YY CU 1
_ 6200 2107          2     7   2   FBAZ1      YY CU 2
_ 6300 2107          3     7   2   FBAZ1      YY CU 3
_ 6400 2107          4     7   2   FBAZ1      YY CU 4
_ 6500 2107          5     7   2   FBAZ1      YY CU 5
_ 6600 2107          6     7   2   FBAZ1      YY CU 6
_ 6700 2107          7     7   2   FBAZ1      YY CU 7
```

Figure 4-35 HCD - Control Unit List: example after I/O Autoconfiguration with values changed

In the View Control Unit Definition panel (Figure 4-36), note the following changes:

- The `Serial-#` now indicates the `Serial number` of the controller.
- Processor SCZP501 is now connected with six static CHPIDs and two dynamic CHPIDs.

```
----------------------- View Control Unit Definition -----------------------
                                                         Row 1 of 7 More:       >
Command ===> ____________________________________________ Scroll ===> CSR

Control unit number . : 6000          YY CU 0
Control unit type . . : 2107          Serial number . . . : FBAZ1

Connected switch.ports: 61.08 61.0C 61.58 61.85 62.08 62.0C 62.58 62.85




ENTER to continue.

           ----------------Channel Path ID . Link Address-----------------
Proc.CSSID 1------ 2------ 3------ 4------ 5------ 6------ 7------ 8------
SCZP201.0  40.08   42.58   41.08   43.58
SCZP201.1  40.08   42.58   41.08   43.58
SCZP401.0  61.08   64.58   62.08   65.58
SCZP401.2  61.08   64.58   62.08   65.58
SCZP401.3  61.08   64.58   62.08   65.58
SCZP501.0  65.85   63.58   64.08   5D.0C   60.0C   51.58   *       *
****************************** Bottom of data ******************************
```

Figure 4-36 HCD - View Control Unit Definition: example after I/O Autoconfiguration showing processor SCZP501 now connected

4.6 I/O Autoconfiguration using HCM

Running the I/O Autoconfiguration process using HCM involves these tasks:

- Defining autoconfiguration policies
- Defining logical partition groups for autoconfiguration
- Defining OS configuration groups for autoconfiguration
- Setting keywords for autoconfiguration policies
- Completing the automatic I/O discovery process
- Applying updates to the autoconfiguration proposals

These procedures assume that you have a working knowledge of HCM.

4.6.1 Defining autoconfiguration policies

To define the autoconfiguration policies, complete the following steps:

1. Click **Utilities** → **Autoconfiguration** (Figure 4-37).

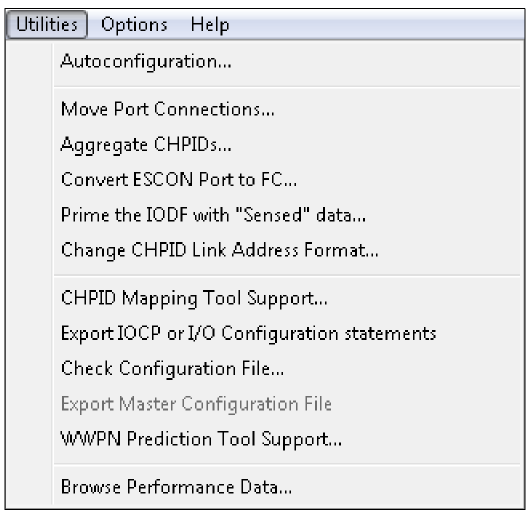

Figure 4-37 HCM - Utilities: Autoconfiguration

2. The first window (Welcome) of the Autoconfiguration wizard opens (Figure 4-38 on page 145). Click **Next**.

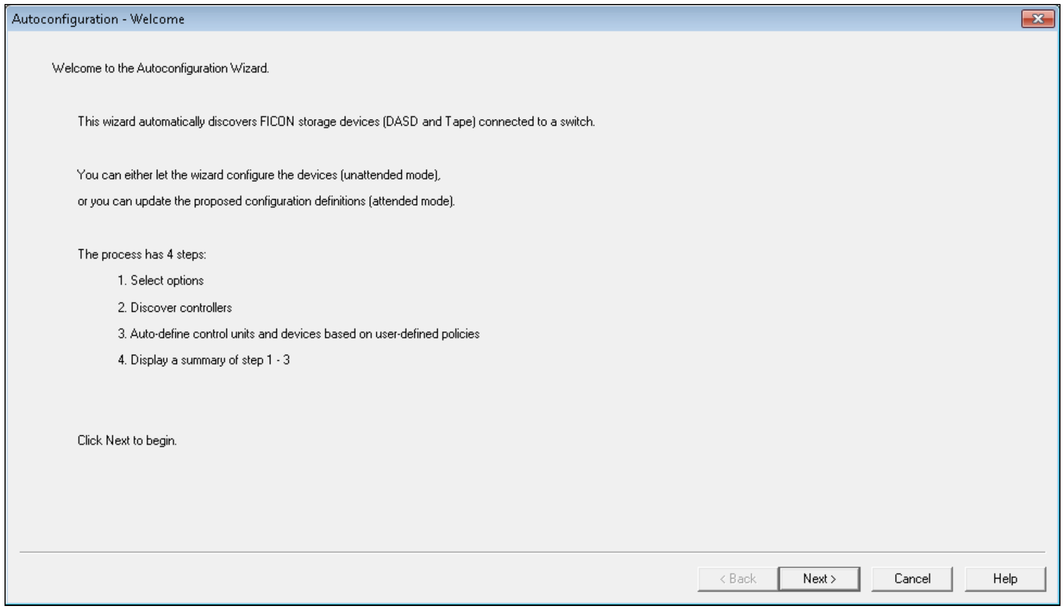

Figure 4-38 HCM - Autoconfiguration: welcome

3. As described in the HCD examples, set up the LP groups and OS groups. Click **Policies** under Specify autoconfiguration policies (Figure 4-39).

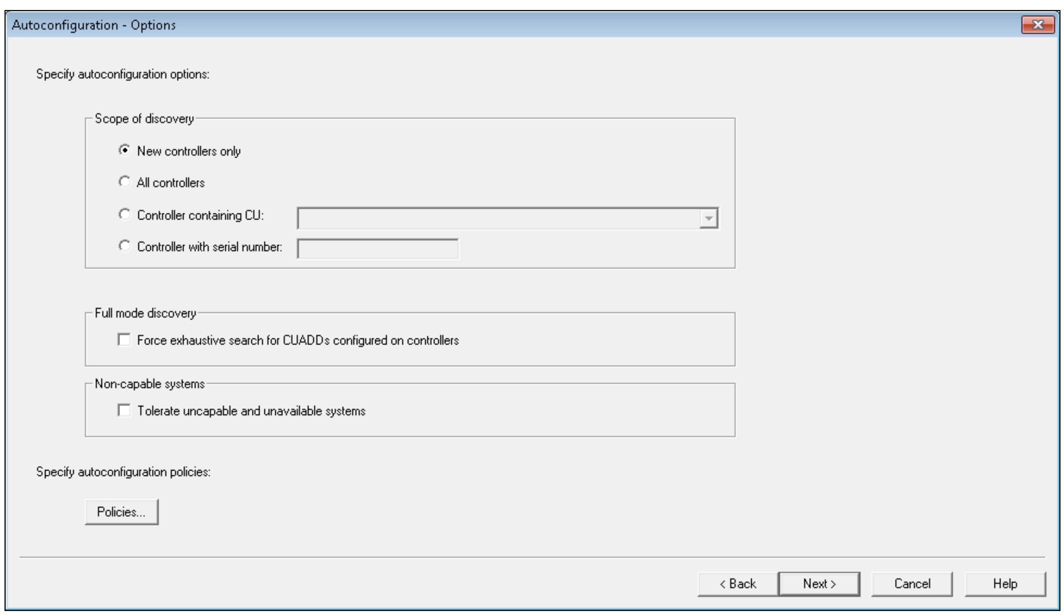

Figure 4-39 HCM - Autoconfiguration: options

The next HCM panel opens (Figure 4-40).

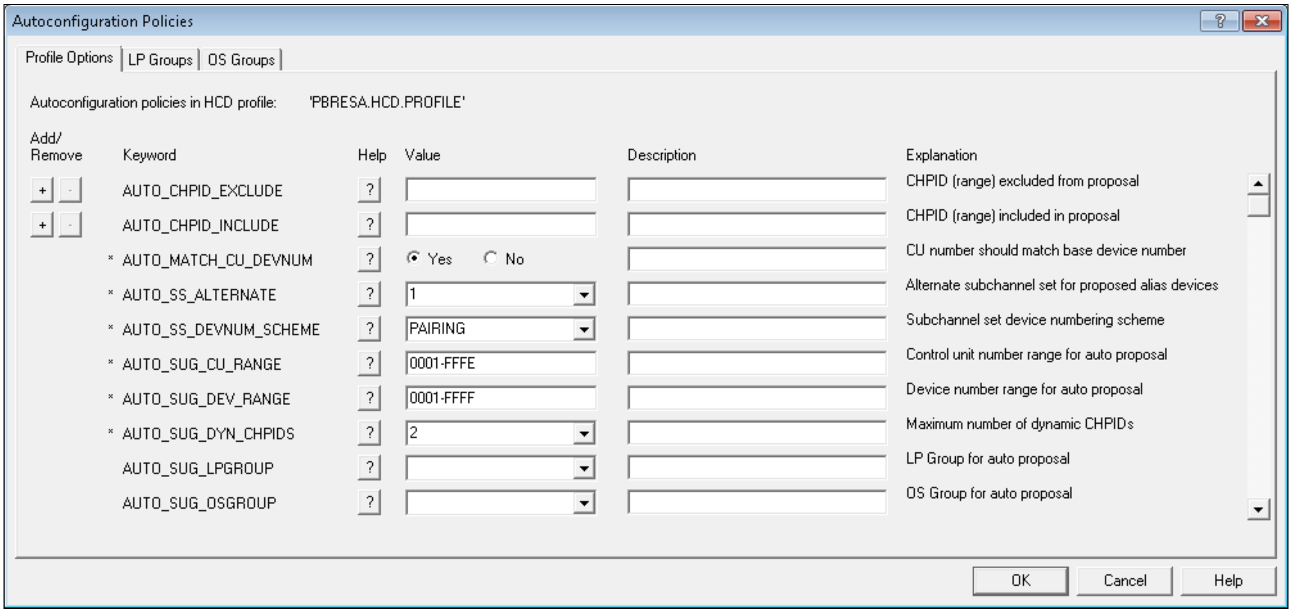

Figure 4-40 HCM - Autoconfiguration Policies: default settings

The default settings are shown in some of the fields for Profile Options:

- If this is the first time you are setting your Profile Options and Policies, the defaults are set and displayed.
- If you previously updated your Profile Options and Policies, they are stored in and retrieved from your `userid.HCD.PROFILE` data set. This example uses the `PBRESA.HCD.PROFILE` data set.

4.6.2 Defining logical partition groups for autoconfiguration

To define logical partition group for autoconfiguration, complete the following steps:

1. Click the **LP Groups** tab (Figure 4-41 on page 147). This window shows that no LP groups are currently defined in this policy. Then, click **Add**.

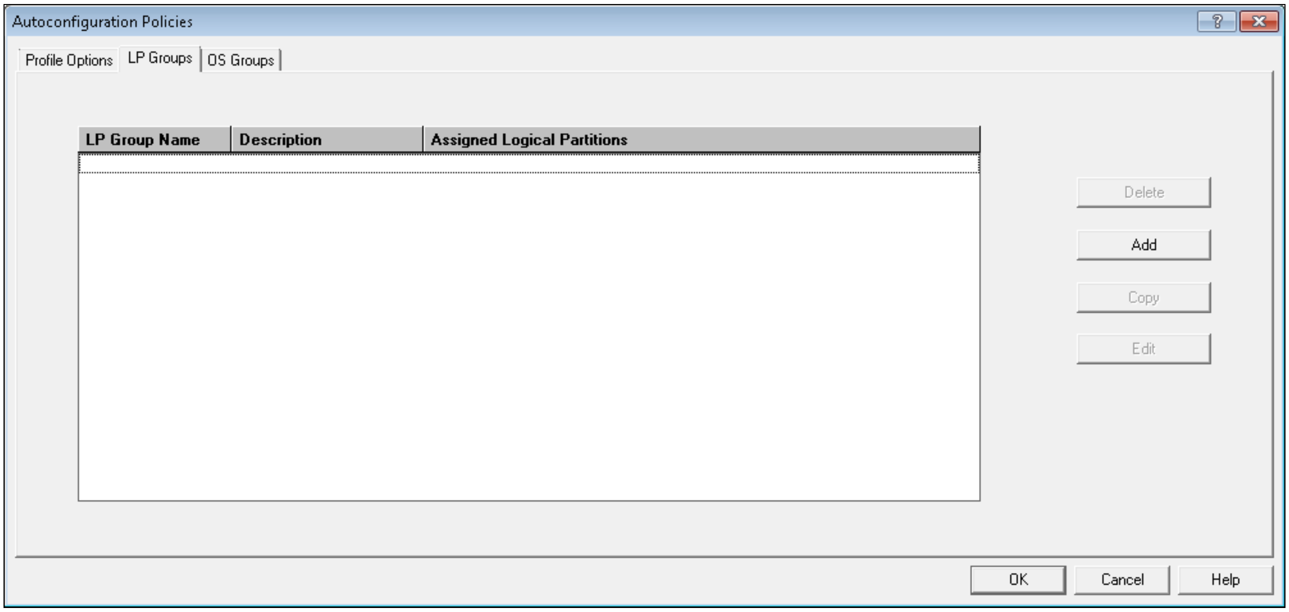

Figure 4-41 HCM - Autoconfiguration Policies: LP Groups with nothing defined

2. In the example, enter an LP Group Name and Description, and then select the **Processor ID** menu to show a list of all defined processors in this IODF (Figure 4-42).

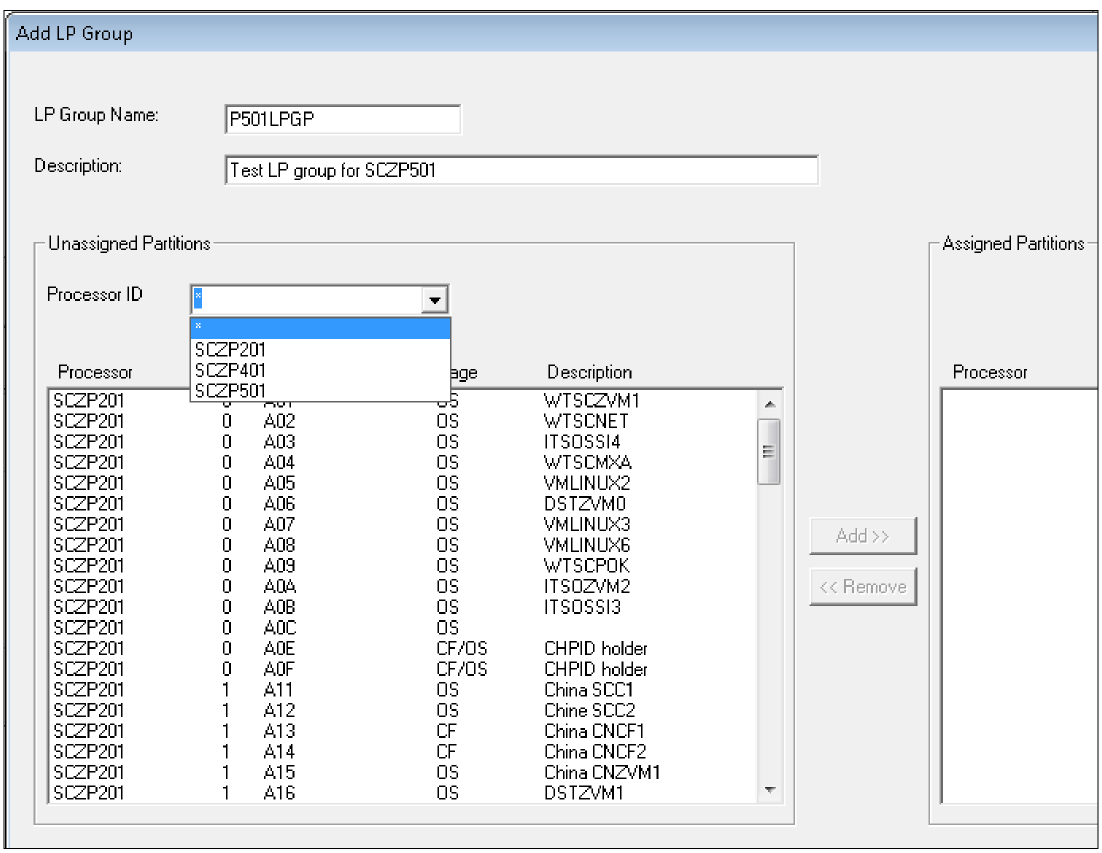

Figure 4-42 HCM - Add LP Group: selecting Processor Ids to add to an LP Group

This option is set to an asterisk (*) by default to show all defined processors and partitions in the box in the lower left corner. You can filter this list by selecting a processor in the Processor ID window. Alternatively, scroll through the list of processors and partitions.

3. Highlight the processor and partition that you want to add to the LP Group and click **Add**. In the example, select processor SCZP501 and partition A01 (Figure 4-43).

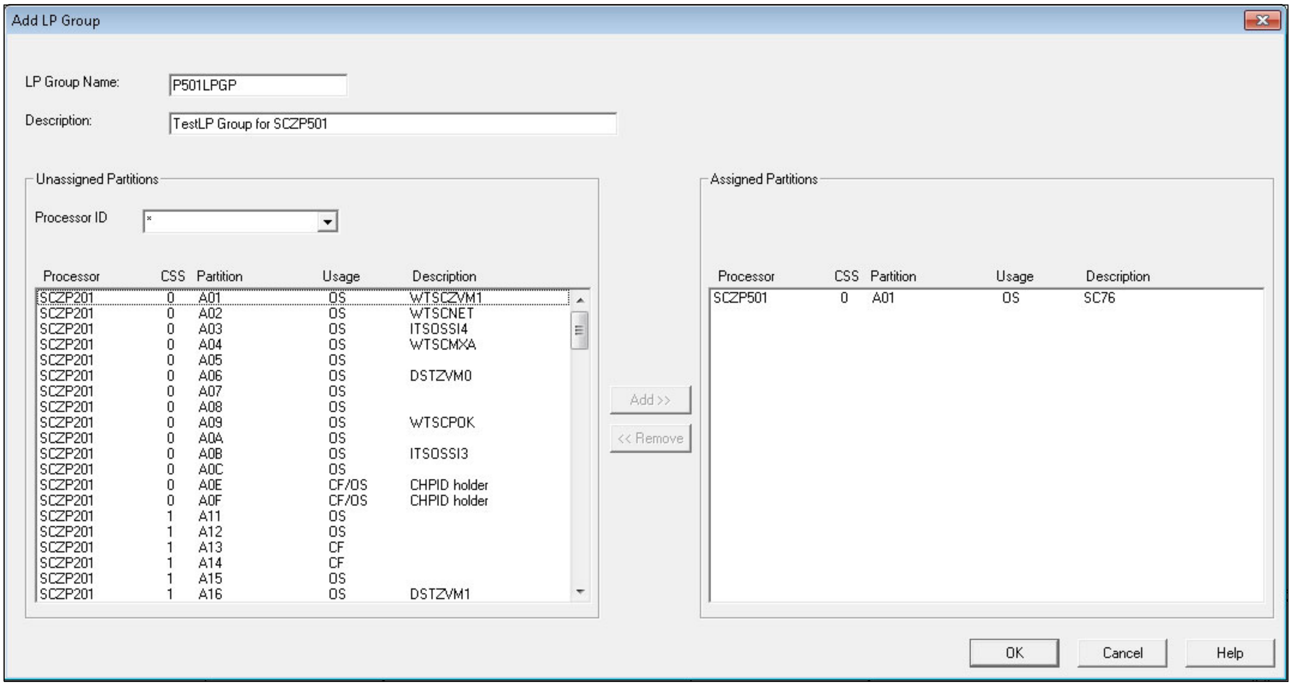

Figure 4-43 HCM - Add LP Group: adding processor ID and partition to an LP Group

4. Click **OK** or add more processors and partitions to the LP Group.
5. Click **OK** (Figure 4-44).

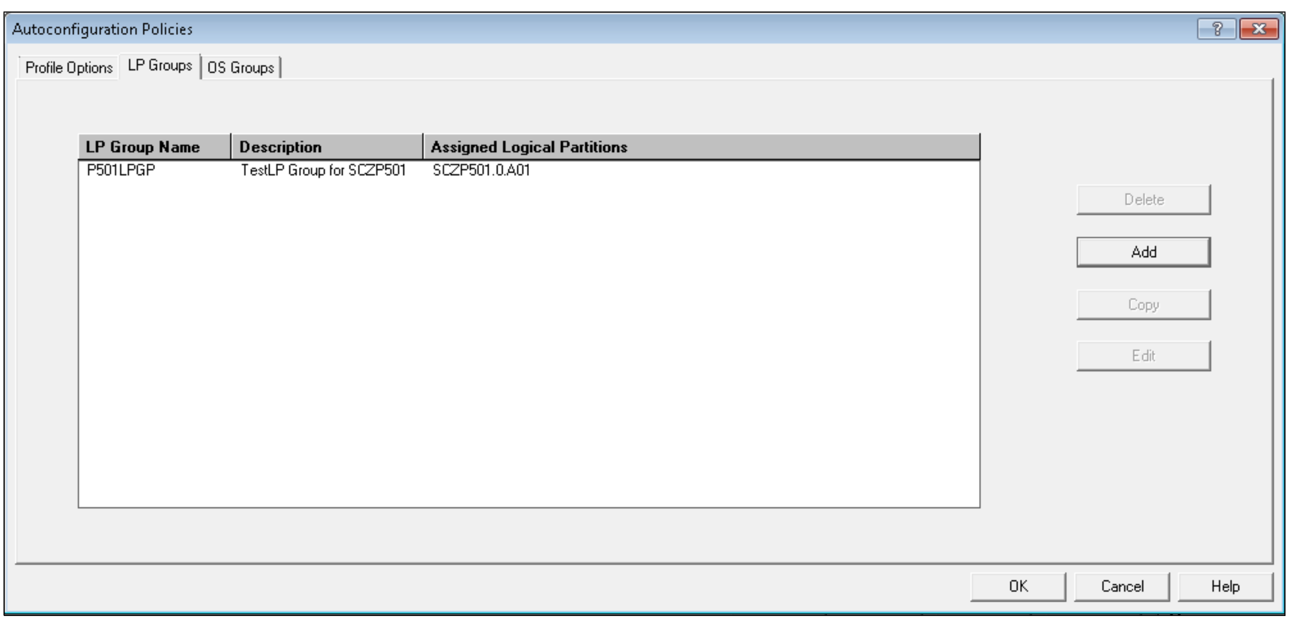

Figure 4-44 HCM - LP Groups: processor ID and partition now added to an LP Group

4.6.3 Defining OS configuration groups for autoconfiguration

To define OS configuration groups for autoconfiguration, complete the following steps:

1. Click the **OS Groups** tab (Figure 4-45). This window shows that there are currently no OS groups defined in this policy. Click **Add**.

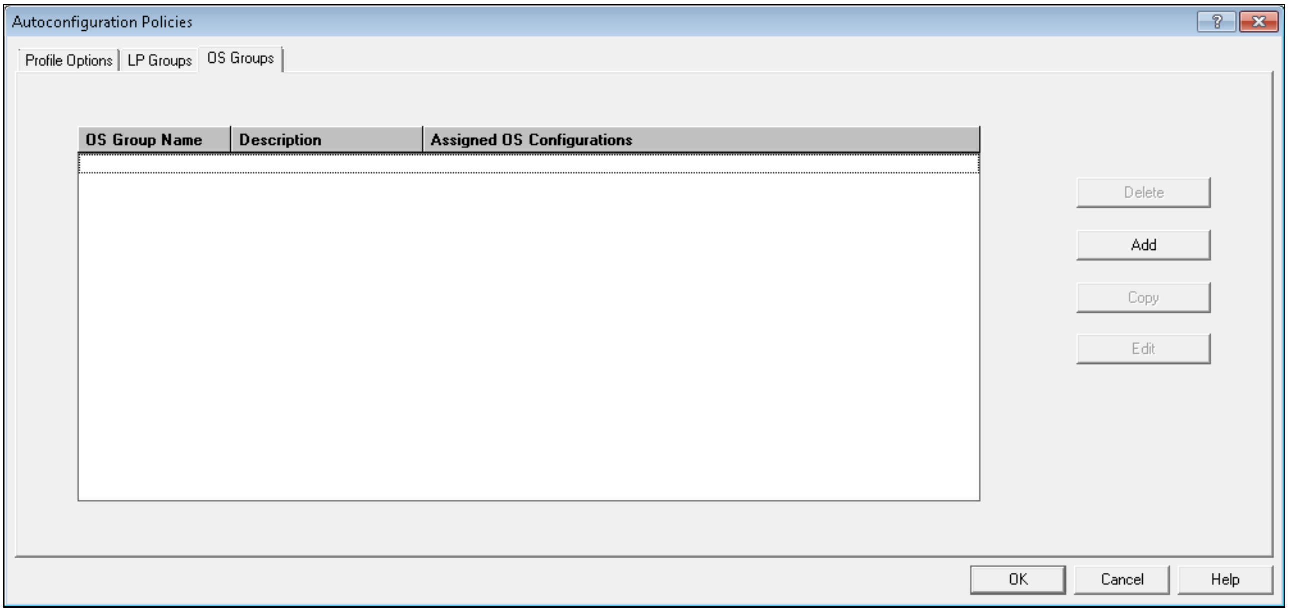

Figure 4-45 HCM - Autoconfiguration Policies: OS Groups with nothing defined

2. In the example, enter an OS Group Name and Description and highlight the OS configuration that you want to add to your OS Group (Figure 4-46).

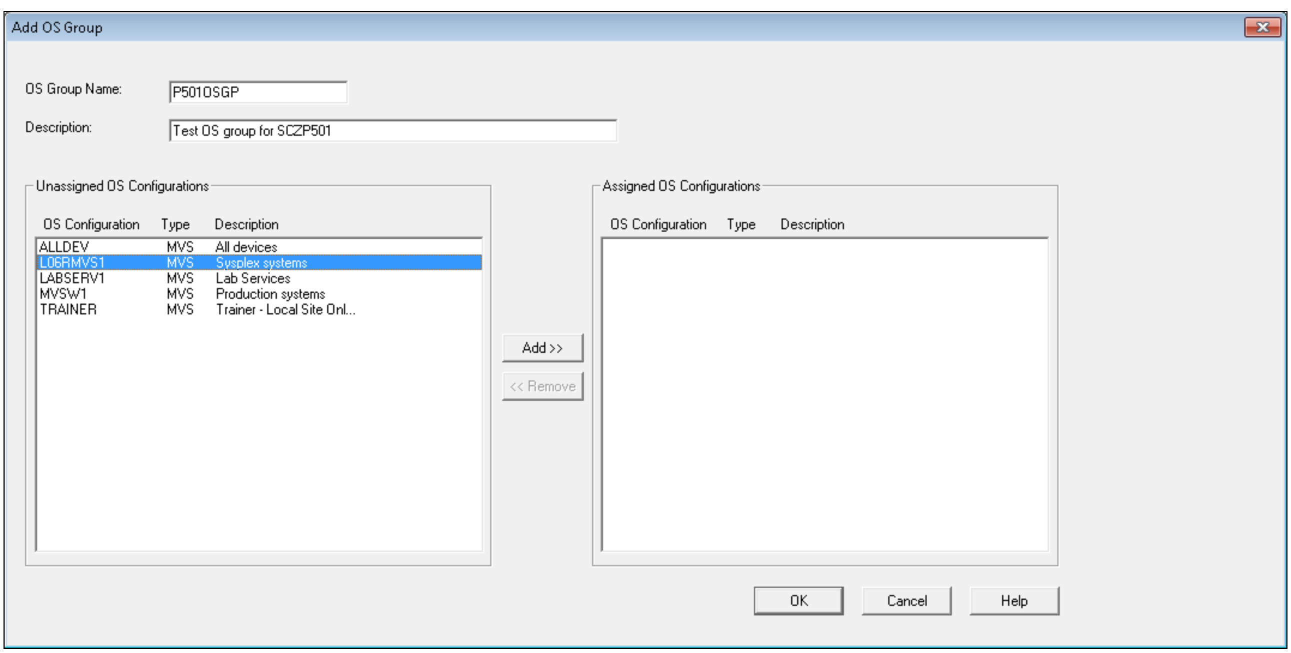

Figure 4-46 HCM - OS Groups: OS configuration list

3. Click **Add** to add the unassigned OS configuration to the assigned OS configuration list and to the OS Group.

4. Click **OK** or add more OS configurations.
5. In the example, add only the L06RMVS1 OS configuration and click **OK** (Figure 4-47).

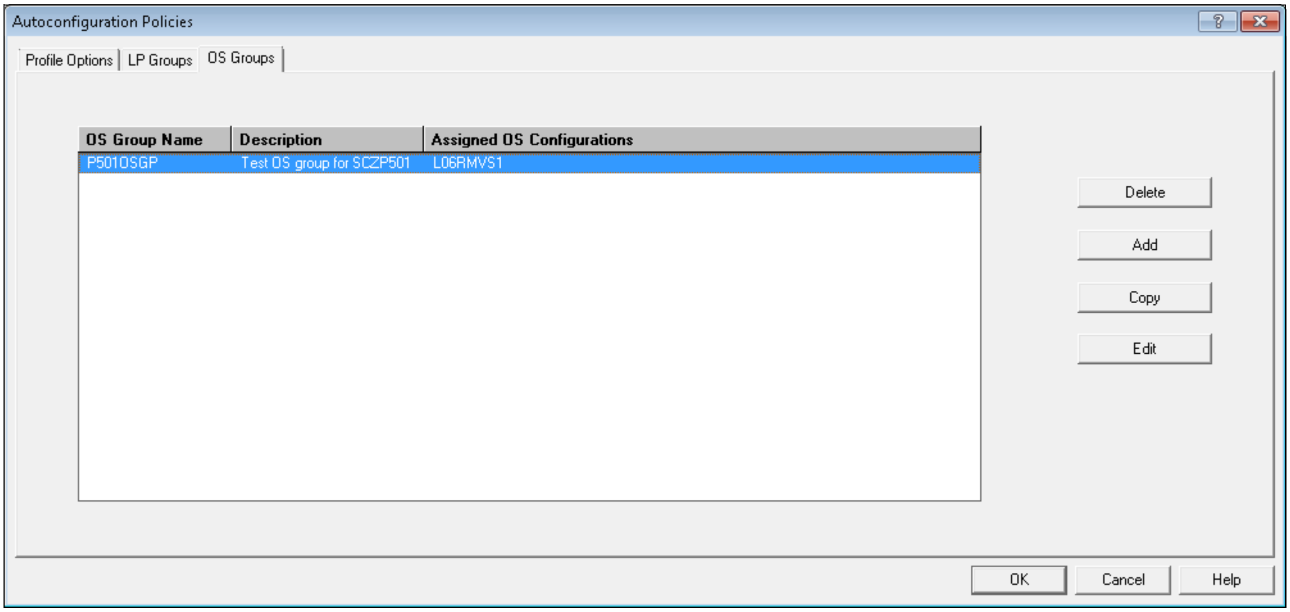

Figure 4-47 HCM - OS Groups: OS configuration added

4.6.4 Setting keywords for autoconfiguration policies

To set keywords for autoconfiguration policies, complete the following steps:

1. Click the **Profile Options** tab (Figure 4-48).

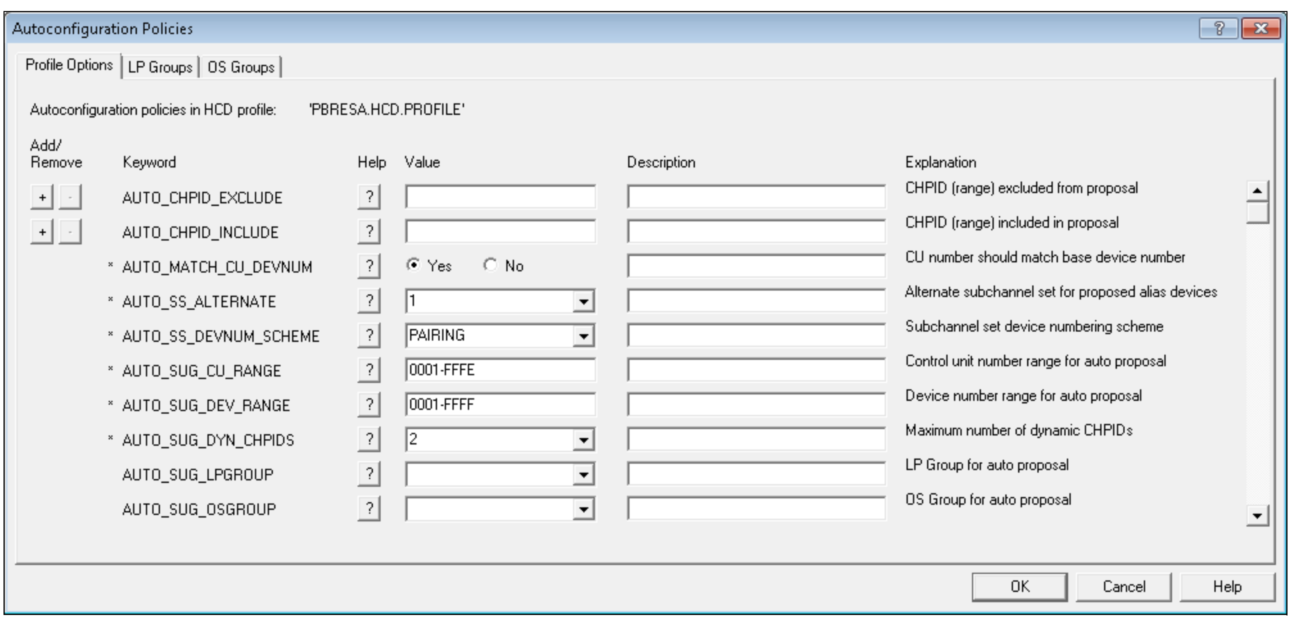

Figure 4-48 HCM - Autoconfiguration Policies: default settings

2. For the example, set the following keywords (Figure 4-49); these values are also found in the pull-down menus next to each data box. Then, click **OK**.
   - AUTO_SS_ALTERNATE: Set to **1** (SS1).
   - AUTO_SUG_DYN_CHPIDS: Set to **2** (dynamic CHPIDs).
   - AUTO_SUG_LPGROUP: Set to **P501LPGP** (LP Group).
   - AUTO_SUG_OSGROUP: Set to **P501OSGP** (OS Group).
   - AUTO_SUG_STAT_CHPIDS: Set to **6** (static CHPIDs).

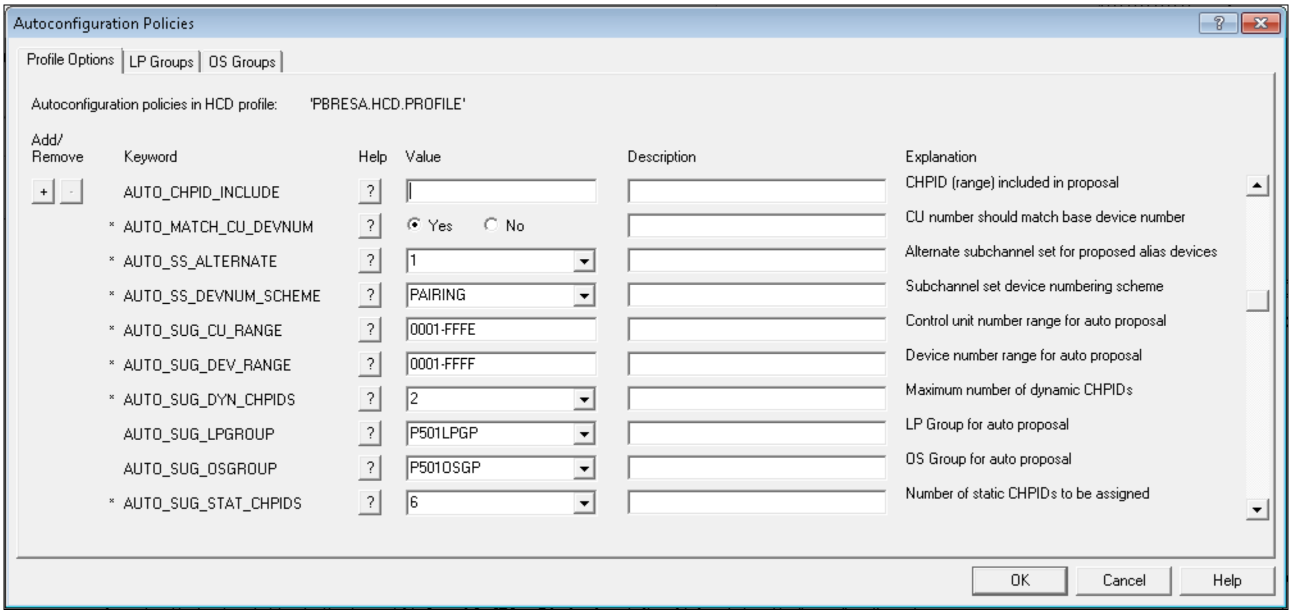

Figure 4-49 HCM - Autoconfiguration Policies: changed from defaults

4.6.5 Completing the automatic I/O discovery process

Now that the profile options and policies are set, start the automatic I/O discovery process. In this example, we search for `New controllers only`.

In the Autoconfiguration - Options window, click **Next** (Figure 4-50 on page 152).

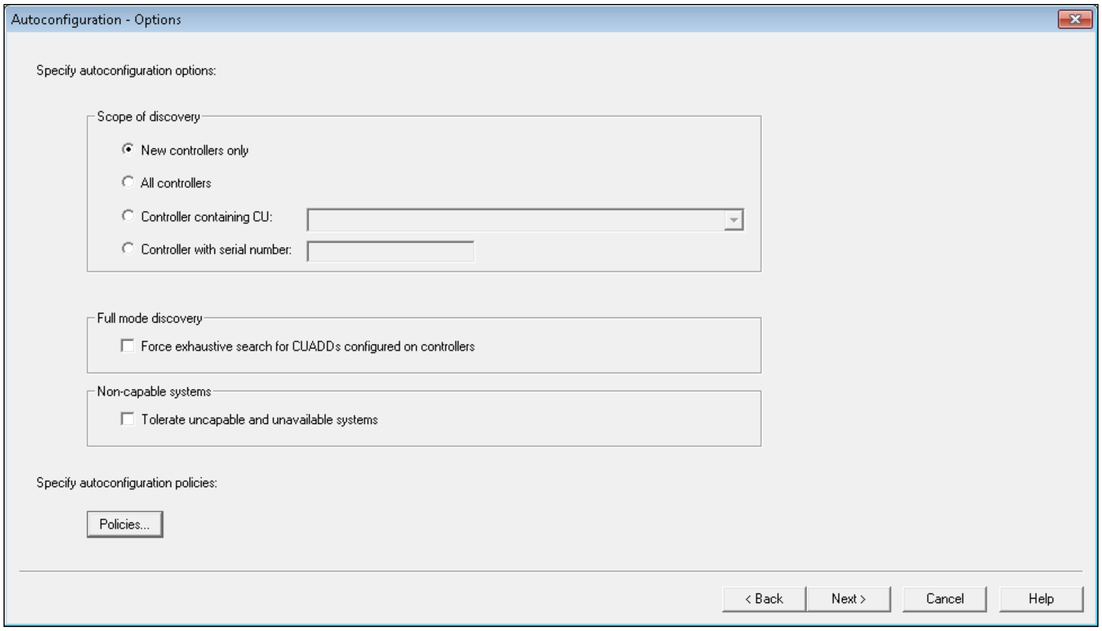

Figure 4-50 HCM - Autoconfiguration Options: beginning the search

The Fabric discovery process starts. When complete, the next window opens (Figure 4-51).

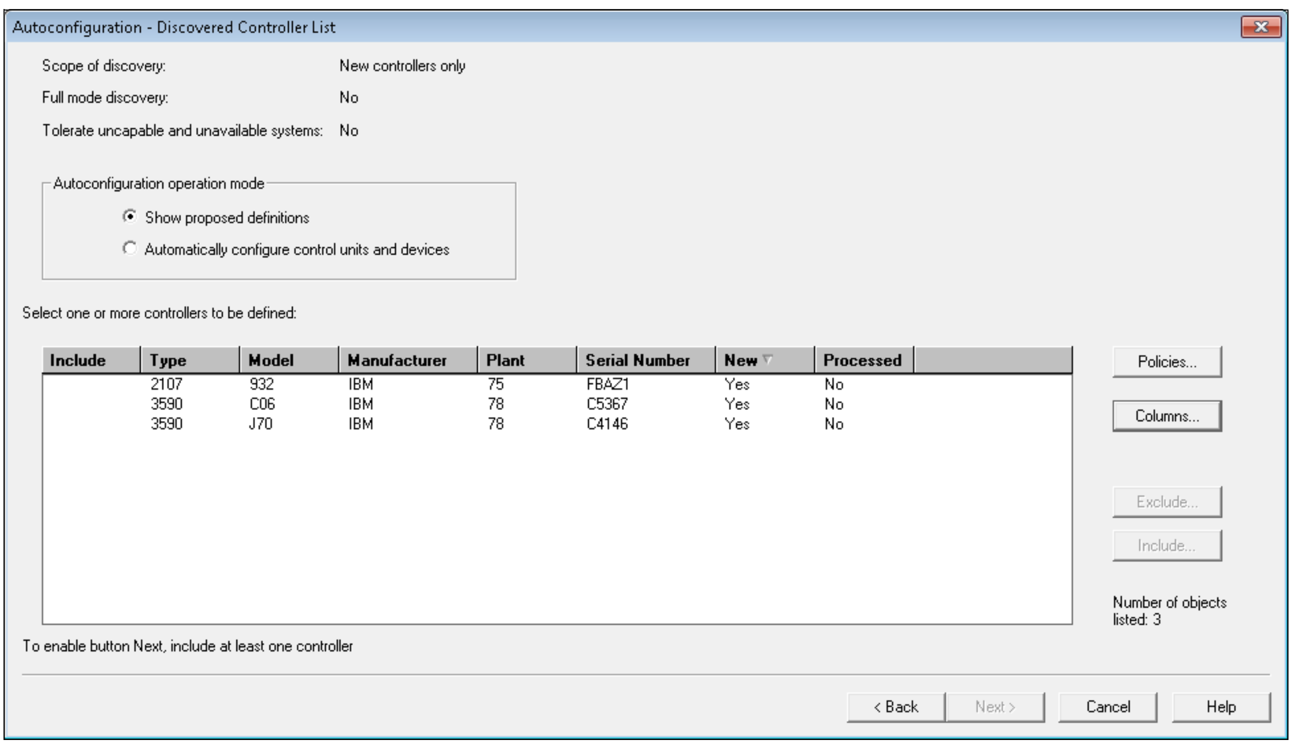

Figure 4-51 HCM - Autoconfiguration: discovered controller list

4.6.6 Applying updates to the autoconfiguration proposals

To apply updates to the autoconfiguration proposals, complete the following steps:

1. Highlight the controller and click **Include** to include that controller for the controller discovery process. Note the check mark in the Include column (Figure 4-52). Click **Next.**

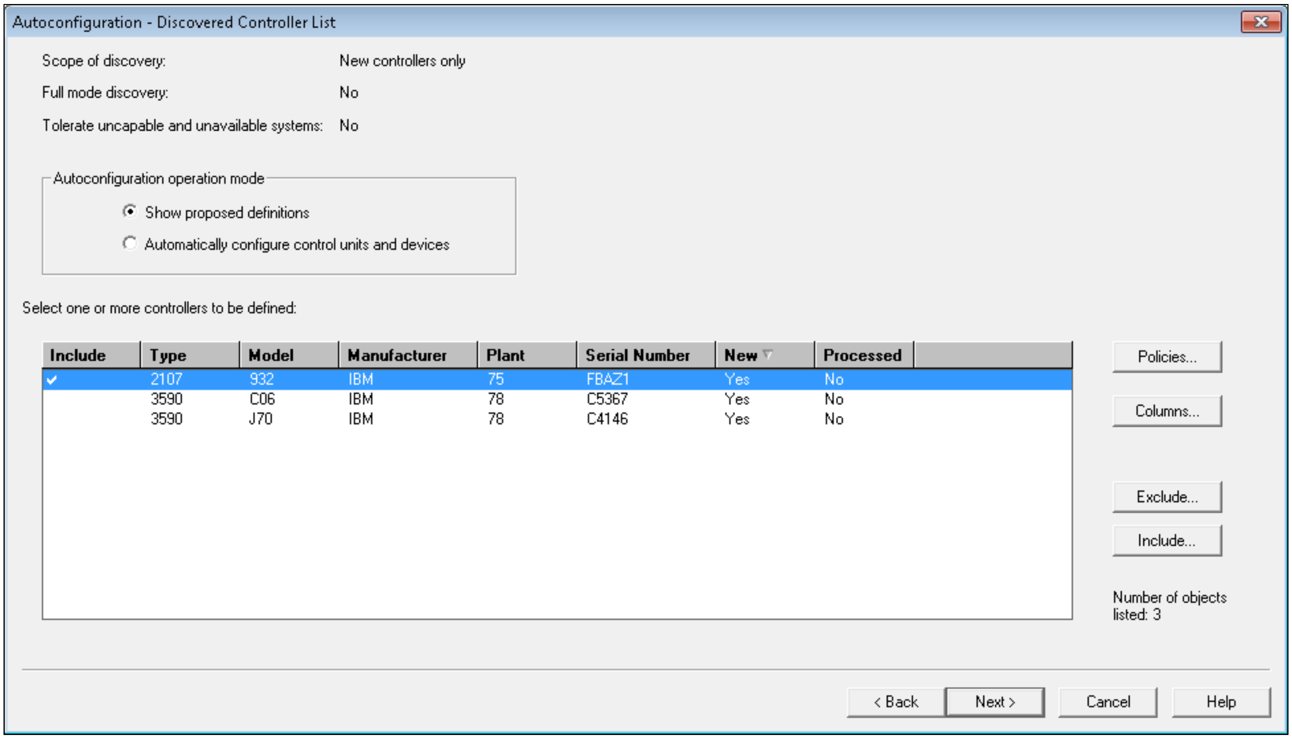

Figure 4-52 HCM - Autoconfiguration: include discovered controller

In the example, HCM issues the CBDG721I message, indicating that the serial number field for the control units has been updated in the IODF with the serial number that was discovered (Figure 4-53).

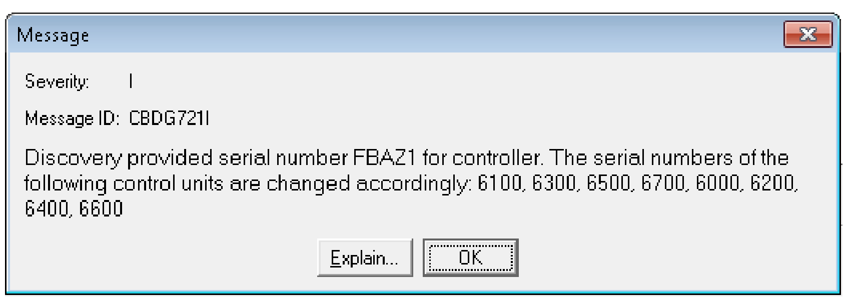

Figure 4-53 HCM - Autoconfiguration: CBDG721I message for serial number changes

Click **OK** to continue.

As a result of the discovery process, the Proposed Control Unit List offers definition proposals for the control units found in the currently processed controller. All eight Control Unit entries are selected as “Include” by default (Figure 4-54 on page 154).

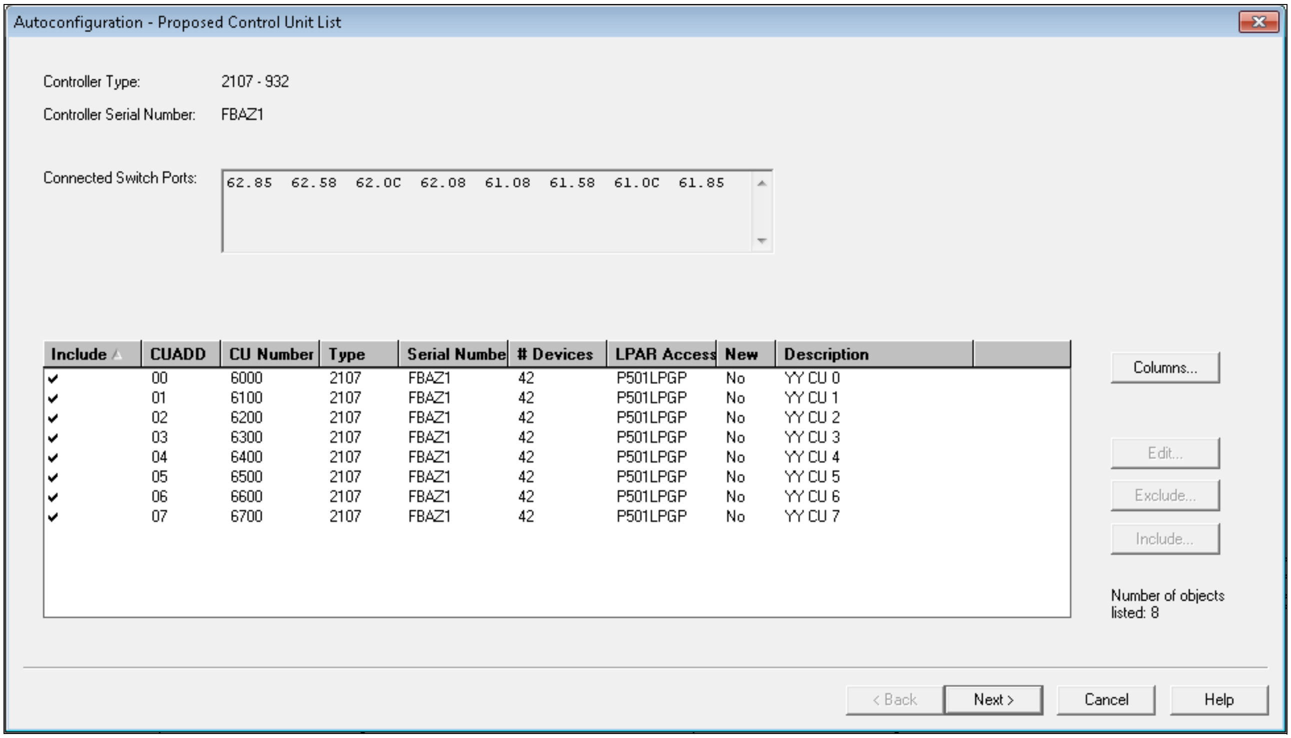

Figure 4-54   HCM - Autoconfiguration: proposed control unit list

2. If you want to edit the proposed control units, highlight the control unit entry and click **Edit**. You can change the Control Unit Number, LPAR Access, and Description. In the example, keep the discovered definitions unchanged (Figure 4-55).

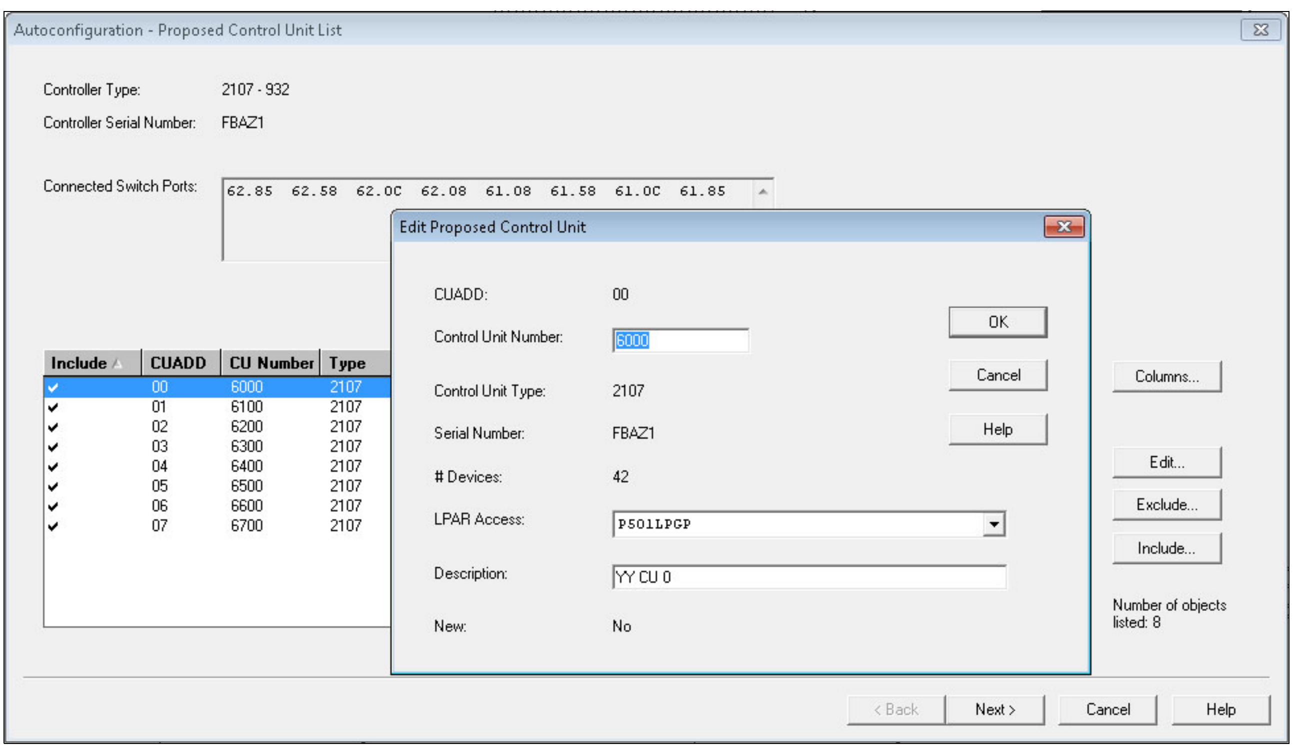

Figure 4-55   HCM - Autoconfiguration: edit proposed control unit list

Click **Next** to continue to the Proposed Control Unit / Device List (Figure 4-56).

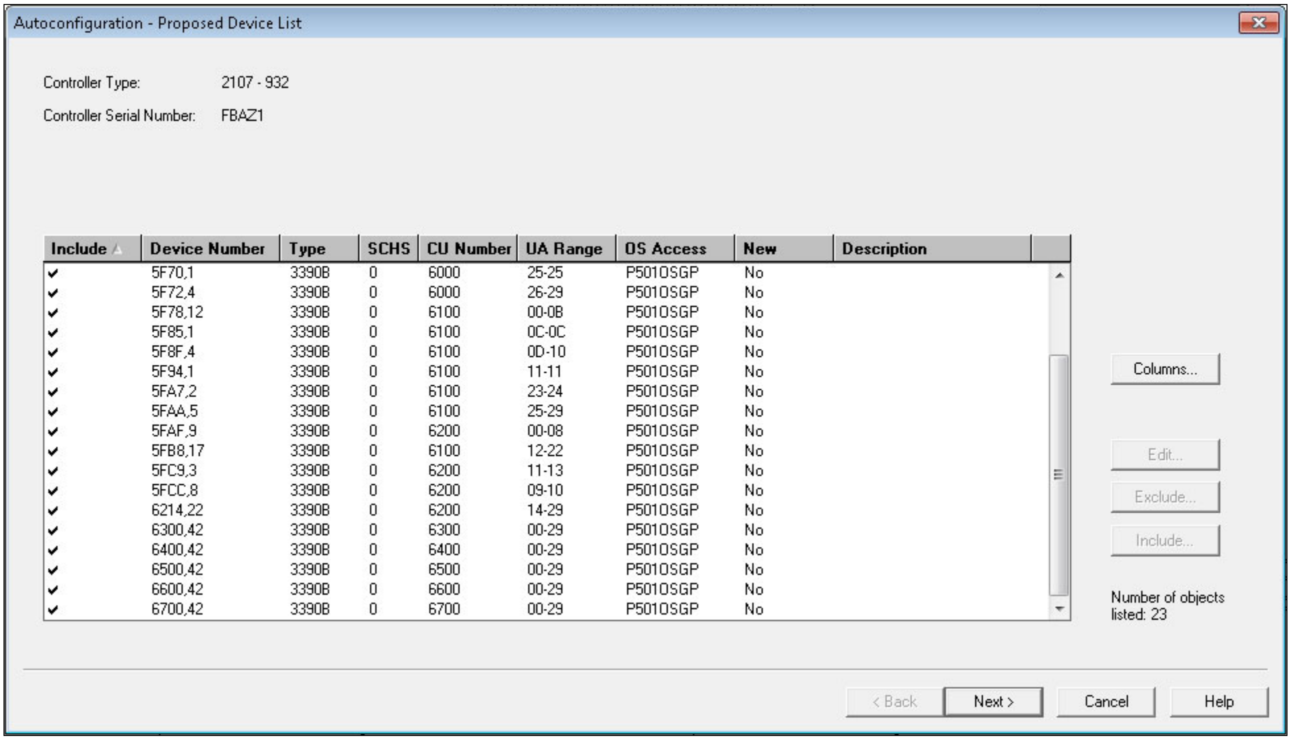

Figure 4-56 HCM - Autoconfiguration: proposed device list

3. Again, all entries are selected as ‘Include’ by default. If there are any device numbers that you do not want to add to the discovery proposal, highlight them and click **Exclude.**
4. In the example, include all the device numbers there were discovered. Click **Next**.
5. HCM issues an Autoconfiguration - Summary report that can be saved to your workstation. Click **Accept** to include these discovered control unit and device definitions into the IODF, or **Reject** if you want to cancel the inclusions (Figure 4-57 on page 156).

   For this example, click **Accept**.

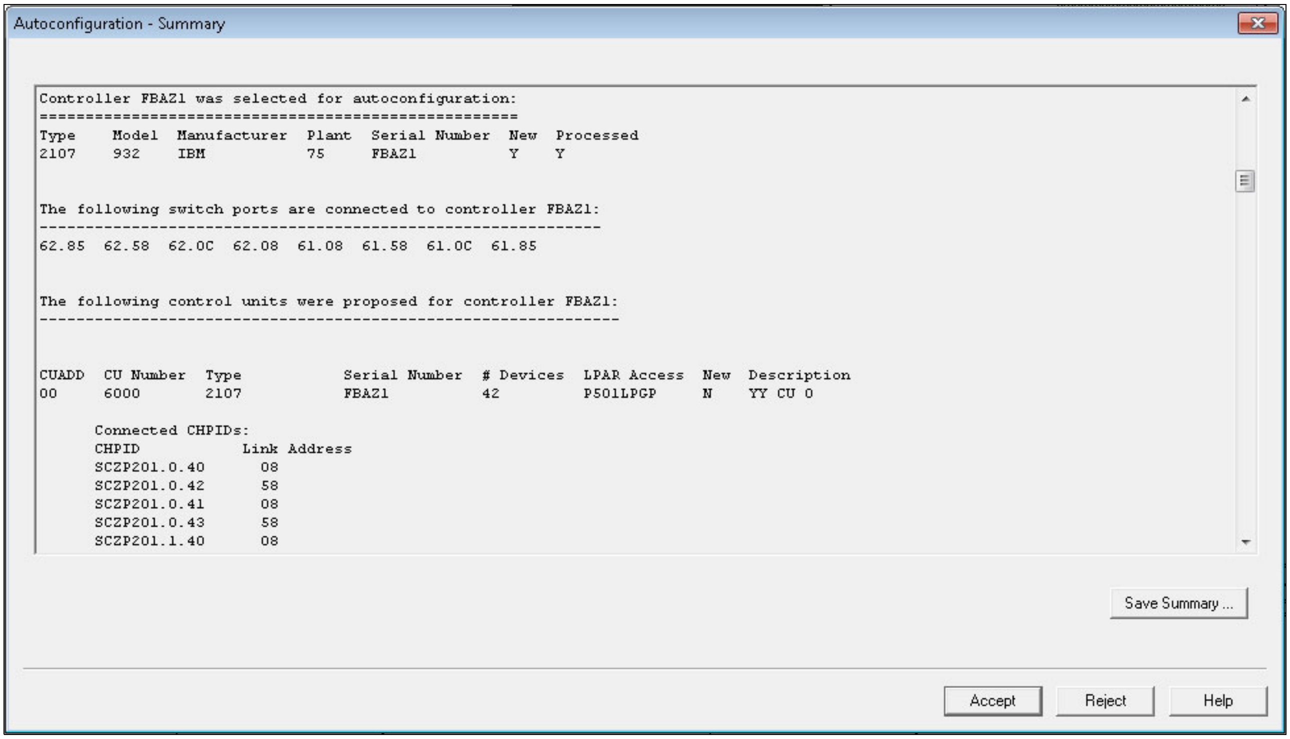

Figure 4-57 HCM - Autoconfiguration: summary report

All discovered and processed controllers now have a `Yes` value in the Processed field (Figure 4-58).

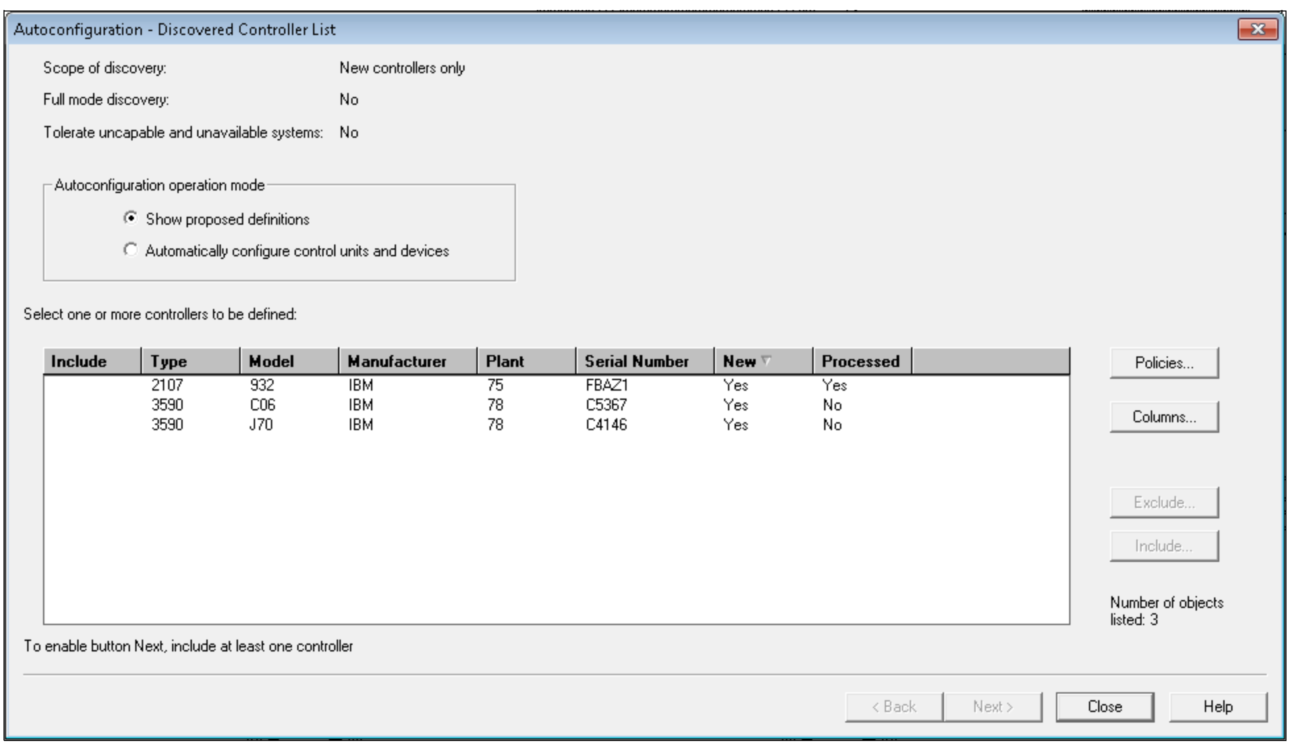

Figure 4-58 HCM - Autoconfiguration: discovered controller processed

6. Click **Close** to complete the Autoconfiguration Wizard process *or* highlight and select **Include** for one or more unprocessed controllers.

Figure 4-59 shows the updated serial number FCAZ1 and discovered paths (six static and two dynamic) added to the control unit definitions for processor SCZP501.

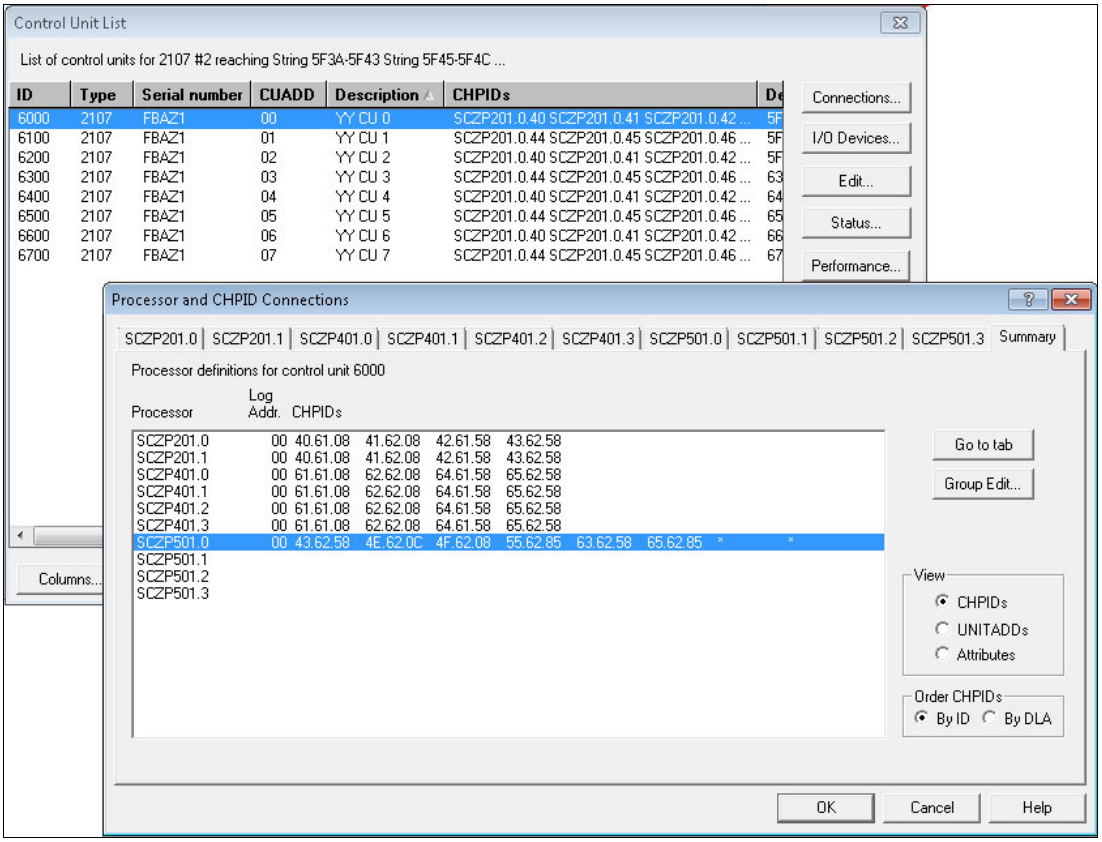

Figure 4-59 HCM - Autoconfiguration: controller with discovered paths added

5

Upgrading from an IBM zEC12 to IBM z13

This chapter describes how to upgrade an existing IBM zEC12 to an IBM z13.

You can also upgrade from a z196 to a z13. However, only examples of upgrading or replacing a zEC12 are shown in detail. An upgrade includes all frames, drawers, and new I/O features.

Because a wide variety of environments exists, the results that are achieved in your environment might differ from those that are described here. Nevertheless, the steps in this chapter can provide enough information for you to replicate the approach in your own environment.

This chapter includes the following topics:

- Scenario overview
- HCD: Migrating the existing 2827 IODF
- OSA: Saving and restoring configuration data
- HCD: Validating the 2964 work IODF
- CMT: Assigning PCHIDs to CHPIDs
- HCD: Updating the 2964 work IODF with PCHIDs
- HCD: Building the 2964 production IODF
- HCD/HMC: Loading the 2964 processor IOCDS
- HMC: Creating Reset Profile for activation

5.1 Scenario overview

This overview describes the upgrade scenario including the configuration process, migration path and planning considerations, and the MES upgrade scenario.

5.1.1 The configuration process

The ten I/O configuration steps that are described in the "I/O configuration process" topic of *I/O Configuration Using z/OS HCD and HCM,* SG24-7804 are used for the example scenario.

Figure 5-1 shows the general process flow that is followed in this example. The numbered steps are described after the figure.

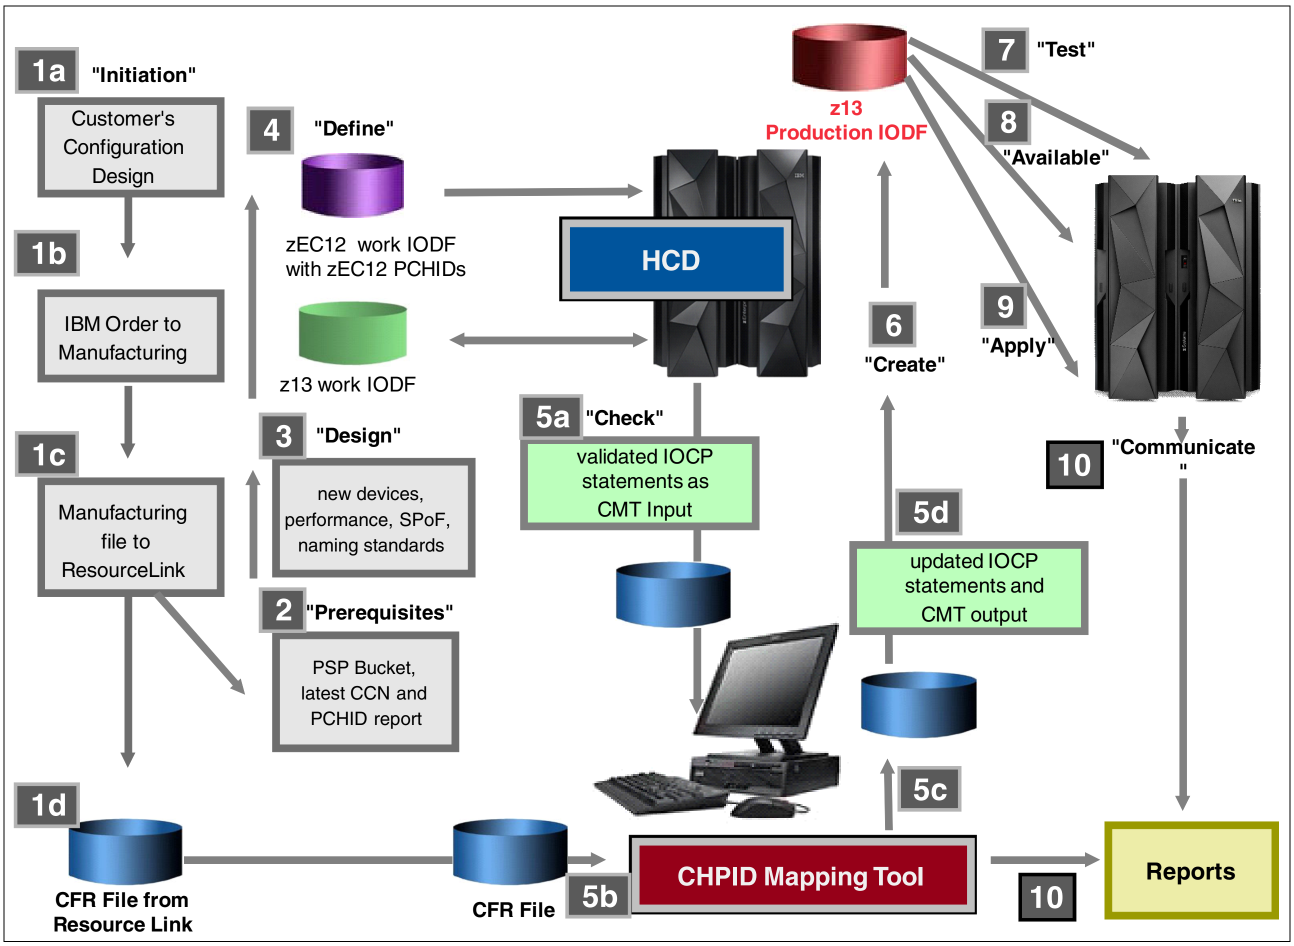

Figure 5-1 Overall configuration process flow

The steps in the figure are as follows:

1. Initiation
   a. When planning to migrate to an IBM z13, the IBM Technical Support team can help you define a configuration design that meets your needs. The configuration is then used during the ordering process.
   b. The IBM order for the configuration is created and passed to the manufacturing process.

c. The manufacturing process creates a configuration file that is stored at the IBM Resource Link website. This configuration file describes the hardware being ordered. This data is available for download by the client installation team.

d. A New Order report is created that shows the configuration summary of what is being ordered along with the Customer Control Number (CCN). The CNN can be used to retrieve `CFReport` (a data file that contains a listing of hardware configuration and changes for a central processor complex (CPC)) from Resource Link.

2. Prerequisites

   Ensure that you have the current PSP Bucket installed. Also, run the SMP/E report with fix category (FIXCAT) exceptions to determine whether any program temporary fixes (PTFs) must be applied. Ensure that you have the most current physical channel ID (PCHID) report and CCN from your IBM System Service Representative.

3. Design

   When you plan your configuration, consider this information:

   – Naming standards
   – FICON switch and port redundancy
   – Adequate I/O paths to your devices for performance
   – OSA channel-path identifier (CHPID) configuration for network and console communications
   – Coupling facility connections internally and to other systems.

   Because the IBM z13 server does not support attachment to the IBM Sysplex Timer, you must consider how the IBM z13 will receive its time source. An IBM z13 cannot join a Coordinated Time Network (CTN) that includes a z10 or earlier as a member. Because the z10 was the last server that supported the IBM Sysplex Timer (9037) connectivity, the z13 cannot be configured as a member of a mixed CTN. The z13 can only join an STP-only CTN.

   When you are planning to replace a z196 or zEC12 with a new z13, plan the replacement of channels that are not supported on IBM z13. You must carefully plan how to replace those, for instance, ISC-3 to HCA3-O or ICA SR for connectivity between z13 and z13.

   You might need to increase CF storage size when you replace z196 or zEC12 with z13. Coupling Facility Control Code (CFCC) level 20 requirements might differ from CFCC level 19 and earlier. Use the CFSizer Tool to get the new CF storage requirements.

4. Define

   The existing z196 or zEC12 I/O configuration is used as a starting point for using Hardware Configuration Definition (HCD). The z196 or zEC12 production input/output definition file (IODF) is used as input to HCD to create a work IODF that becomes the base of the new IBM z13 configuration.

   When the new IBM z13 configuration is added and the obsolete hardware is deleted, a validated version of the configuration is saved in an IBM z13 validated work IODF.

5. Check

   a. From the validated work IODF, create a file that contains the IBM z13 IOCP statements. This IOCP statements file is transferred to the workstation used for the CHPID Mapping Tool (CMT). Hardware Configuration Manager (HCM) can also be used here to transfer the IOCP deck to and from the CMT.

   b. The configuration file that is created by the IBM Manufacturing process in step 1d is downloaded from Resource Link to the CMT workstation.

The CHPID Mapping Tool (CMT) uses the input data from the files to map logical channels to physical ones on the new IBM z13 hardware.

You might make decisions in response to the following situations, among others:

i. Resolving situations in which the limitations on the purchased hardware cause a single point of failure (SPoF). Consider purchasing more hardware to resolve some SPoF situations.

ii. Prioritizing certain hardware items over others.

c. After the CMT processing finishes, the IOCP statements contain the physical channels to logical channels assignment that is based on the actual purchased hardware configuration.

The CHPID Mapping Tool (CMT) also creates configuration reports to be used by the IBM System Service Representative and the installation team.

The file that contains the updated IOCP statements created by the CMT, which now contains the physical channels assignment, is transferred to the host system.

d. Use HCD, the validated work IODF file created in step 5a, and the IOCP statements updated by the CMT to apply the physical channel assignments created by the CMT to the configuration data in the work IODF.

6. Create

After the physical channel data is migrated into the work IODF, an IBM z13 production IODF is created and the final IOCP statements can be generated. The installation team uses the configuration data from the IBM z13 production IODF when the final power-on reset is done, yielding an IBM z13 with an I/O configuration ready to be used.

7. Test

IODFs that are modifying existing configurations can be tested in most cases to verify that the IODF is making the intended changes.

8. Available

a. If you are upgrading an existing z196 or zEC12, you might be able to use HCD to write an IOCDS to your system in preparation for the upgrade. If you can write an IOCDS to your current system in preparation for upgrade, do so and let the IBM System Service Representative know which IOCDS to use.

> **Tip:** Using the HCD option “Write IOCDS” in preparation of an upgrade is the preferred method for writing the initial IOCDS when upgrading from a z196 or a zEC12 to an IBM z13. This scenario uses the HCD option *Write IOCDS process*.

b. If the z196 or zEC12 is not network connected to the CPC where HCD is running, or if you are not upgrading or cannot write an IOCDS in preparation for the upgrade, use HCD to produce an IOCP input file. Download this input file to a USB flash drive.

9. Apply

The new production IODF can be applied to the IBM z13 in these ways:

– Using the Power-on Reset process
– Using the Dynamic IODF Activate process

10. Communicate

Communicating new and changed configurations to operations and the appropriate users and departments is important.

5.1.2 Migration path considerations

The migration path from a z196 or zEC12 to an IBM z13 can be either in the form of a field upgrade to the existing z196 or zEC12, or a replacement (push/pull) of an existing z196 or zEC12 with a new IBM z13. Note the following points:

- A *field upgrade* means that the existing z196 or zEC12 processor serial number is retained during the upgrade.
- A *replacement* of the existing z196 or zEC12 by a new IBM z13 implies physically removing (push) the z196 or zEC12 and bringing in a new IBM z13 (pull) to take its place. The replacement IBM z13 has a new serial number that is different from that of the existing z196 or zEC12.

This chapter documents a *field upgrade* scenario.

5.1.3 Planning considerations

The following I/O features (listed with feature codes, FC) can be ordered for a new z13:

- FICON Express16S LX (long wavelength - 10km) - FC 0418
- FICON Express16S SX (short wavelength) - FC 0419
- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- RoCE-Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Crypto Express5S - FC 0890
- Flash Express - FC 0403
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171
- PCIe-O SR Integrated Coupling Adapter - FC 0172

The following features, if present in a z196 or zEC12, can be carried forward when you upgrade to an IBM z13:

- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- FICON Express8 LX (long wavelength - 10 km) - FC 3325
- FICON Express8 SX (short wavelength) - FC 3326
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- OSA-Express4S GbE LX (long wavelength) - FC 0404
- OSA-Express4S GbE SX (short wavelength) - FC 0405
- OSA-Express4S 10 GbE LR (long reach) - FC 0406
- OSA-Express4S 10 GbE SR (short reach) - FC 0407
- OSA-Express4S 1000BASE-T Ethernet - FC 0408
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Flash Express - FC 0402

- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171

The following features are not supported on IBM z13:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- OSA-Express2 10 GbE Long Reach
- OSA-Express
- Crypto Express3
- Crypto Express2
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode
- ICB-2
- ICB-3
- ICB-4
- PCIXCC and PCICA
- Parallel channels (use FICON and ESCON converters)

Table 5-1 lists the channel types as described in an I/O configuration data set (IOCDS) that are supported by the z196 (2817), zEC12 (2827), and z13 (2964).

Table 5-1 Channels, links, and adapters with CHPID type and support

Channels	CHPID type	2817 support	2827 support	2964 support
ESCON channels: Connection Channel (ESCON architecture) Channel to Channel (connects to CNC) Converter Channel Path (for BL types) Converter Channel Path (for BY types)	CNC CTC CVC CBY	Up to 240 Up to 240 Up to 240 Up to 240	- No - -	- No - -
FICON native channels that attach to FICON directors or directly to FICON control units: FICON Express 4 SX and LX FICON Express 8 SX and LX FICON Express 8S SX and LX FICON Express 16S SX and LX	FC FC FC FC	Carry forward up to 288 Yes No	Carry forward Carry forward up to 320 No	No up to 64 Carry Forward up to 320
FICON channels that attach to Fibre Channel devices, switches, directors, or Fibre-Channel-to-SCSI bridges	FCP	Yes	Yes	Yes
ISC-3 peer mode channels (connects to another CFP)	CFP	Up to 48	Carry forward	No
IC peer channels (connects to another ICP)	ICP	Up to 32	Up to 32	Up to 32
HCA2-O 12x InfiniBand host channel adapters **HCA2-O 1x LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 32	Up to 32 links Up to 32 links	No No

Channels	CHPID type	2817 support	2827 support	2964 support
HCA3-O 12x InfiniBand host channel adapters **HCA3-O LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 48	Up to 32 links Up to 64 links	Up to 32 links Up to 64 links
PCIe-O SR Integrated Coupling Adapter	CS5	No	No	Up to 32 links
HiperSockets (IQDIO) channels	IQD	Up to 32	Up to 32	Up to 32
OSA- Express2 GbE LX/SX	OSD and OSN	Up to 48 ports carried forward	No	No
OSA- Express2 1000BASE-T	OSE, OSD, OSC, and OSN	Yes	No	No
OSA- Express2 10 GbE LR	OSD	Yes	No	No
OSA-Express3 GbE LX/SX	OSD and OSN	Up to 96 ports	Carry forward	No
OSA-Express3 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	OSC, OSD, OSE, and OSN: Yes OSM: No	Carry forward	No
OSA- Express3 10 GbE LR/SR	OSD and OSX	OSD: Yes OSX: No	Carry forward	No
OSA-Express4S GbE LX/SX	OSD	Up to 96 ports	Up to 96 ports	Carry forward
OSA-Express4S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Carry forward
OSA-Express4S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Carry forward
OSA-Express5S GbE LX/SX	OSD	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Up to 48 ports
10GbE RoCE Express SR	ROCE	No	Up to 32 ports	Up to 32 ports
zEDC Express Enterprise Data Compression	ZEDC	No	Max 8 features	Max 8 features

Consider the following information when you plan your configuration:

- Coupling links
- Hardware Management Console (HMC)
- Software support
- Open Systems Adapter - Integrated Console Controller
- Fibre Channel Protocol
- CPC name versus Processor ID
- Local system name

Coupling links

Only the following coupling facility CHPIDs are supported:

- CHPID type CIB: PSIFB links connecting to an HCA3-O 12x and HCA3-O LR 1x
- CHPID type ICP: Internal Coupling links.
- CHPID type CS5: Integrated Coupling adapter PCIe-O SR

Considerations: Coupling links can be defined as both Coupling and STP links, or STP-only links. IBM z13 does not support the ISC-3 features.

Hardware Management Console (HMC)

The HMC can appear either as the current HMC does, or as an HMC that can run code to manage an Ensemble. The current HMC is used to manage, monitor, and operate one or more IBM z Systems servers and their associated logical partitions. An HMC that has ensemble code running is an HMC attached to one or more zEnterprise Systems configured as ensemble members. A particular ensemble is managed by a pair of HMCs in primary and alternate roles.

The HMC has a global (ensemble) management function, whereas the Support Element (SE) has local node management responsibility. When tasks are performed on the HMC, the commands are sent to one or more SEs, which then issue commands to their CPCs.

The z13 requires HMC Application V2.13.0 (driver level 22) or later, and uses only Ethernet for its network connection. The HMC and the SEs do not contain a floppy disk drive. Therefore, you must use a USB flash memory drive to input and back up client configuration data.

Software support

HCD V2.1 or later with the Preventive Service Planning (PSP) bucket for 2964DEVICE and PTFs) is required to define and support some of the new features of the z13.

Open Systems Adapter - Integrated Console Controller

Because support is withdrawn for the 2074 console controllers, consider using OSA-Express4S 1000BASE-T or OSA-Express5S1000BASE-T CHPIDs defined as TYPE=OSC. With this OSA card, you can set up console function that is supported by a configuration file defined on the Support Element for that processor.

Fibre Channel Protocol

When you use CHPIDs defined as TYPE=FCP, consider N-Port ID Virtualization (NPIV).

For more information about FCP CHPIDs and the WWPN prediction tool to manage them, see “Batch Network Analyzer and Compression” on page 47.

CPC name versus Processor ID

HCD allows you to define different processors (logical) to the same CPC (physical). The Processor ID must be unique within the same IODF, but the CPC name does not. Therefore, the CPC name does not need to match the Processor ID. This advantage is useful when you have several processor/logical partition/control unit setups that share a physical CPC within the same IODF. Furthermore, the Processor ID is what is defined for the optional HWNAME parameter in the LOAD member in SYS1.IPLPARM.

The CPC name is coded in HCD option 1.3 under View Processor Definition in the CPC name field under SNA address, along with a Network name. It is the CPC name, and not the Processor ID, that is displayed on the HMC.

When you view the network information for a CPC through the HMC, the SNA address is made up of a network name and CPC name separated by a dot (for example, `USIBMSC.SCZP501`). These values are defined in the Support Element for the CPC. They must match the values that are set in the IODF so that HCD option 2.11 can find the CPC to write an IOCDS in the System z Cluster List.

Local system name

An extra system name, LSYSTEM, is used to identify the local system name of a server when you define PSIFB type=CIB or ICA type=CS5 coupling links.

This data field can be found when you change a CIB-capable processor under HCD option 1.3.

The LSYSTEM field can be set or changed to any one to eight alphanumeric characters. Also, it can begin with either an alphabetic or numeric character. All characters are uppercase.

The following rules determine whether, and where, HCD sets the LSYSTEM keyword automatically:

- If a CIB-capable processor is defined and the CPC name is set but the local system name is not set, HCD sets the local system name to the CPC name.
- If a CIB-capable processor that has not yet defined a CPC name is changed to obtain a CPC name but no local system name, HCD sets the CPC name to the local system name.
- If a non-CIB capable processor is changed to a support level that is CIB capable, and the processor has a CPC name set but no local system name, the local system name defaults to the CPC name.
- If the processor definition is changed such that the local system name is explicitly removed, HCD does not do any defaulting.
- If a processor has a local system name set (whether it has a CPC name or not), any change to the local system name must be done explicitly. There is no implicit name if the CPC name or the support level is changed.
- During Build Production IODF, verification determines whether a local system name is set for the processor if the processor has a CIB channel path defined. If this verification fails, an error message is displayed and the production IODF is not built.

Generally, set the local system name the same as the CPC name.

The following additional keywords are for the ID statement in an IOCP deck:

CPATH Specifies the CSSID and CHPID on the connecting system.

CSYSTEM Connecting System name (LSYSTEM in processor definition).

AID Host Card Adapter ID (found in PCHID report).

PORT Host Card Adapter port.

5.1.4 Miscellaneous equipment specification (MES) upgrade scenario

This scenario describes the configuration steps to upgrade an existing 2827 processor to a 2964 processor. Key factors are as follows:

- HCD requires a new Processor ID for the 2964.
- Generally, keep the same CPC name for the 2964.
- The 2964 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 2827 production IODF.
- The target IODF is a new 2964 work IODF.
- 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.

The example uses a 2827-H43 with a Processor ID of SCZP401 and four CSSs (CSS ID=0, CSS ID=1, CSS=2, and CSS ID=3). This system is replaced with a 2964-N63 with a Processor ID of SCZP501 and six CSSs.

The CPC name SCZP401 and serial number 02-0B8D7 are not changed.

The following CHPID types are migrated:

- OSD, OSC, OSM, and OSX
- FC and FCP
- CIB and ICP
- IQD

The following hardware or CHPID types are *not* supported and *not* migrated to the 2964:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- Crypto Express3
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode

Table 5-2 summarizes the migration options and tool requirements. The process steps are described in 5.2, "HCD: Migrating the existing 2827 IODF" on page 169.

Table 5-2 2827 I/O configuration migrated to a 2964

2827 to 2964	Upgrade existing 2827 to a 2964 (MES upgrade)
Processor ID	Required to change the Processor ID to a new ID
CPC name	Generally should be the same name
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 (see step 5 on page 169)
CHPID Mapping Tool Program	Optional, but quite useful
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

5.2 HCD: Migrating the existing 2827 IODF

The following steps explain how to define an existing 2827 I/O configuration to the new 2964 server by using HCD. You then migrate the channel subsystem and logical partitions from the 2827 to the 2964. Using HCD, the sequence of operations is as follows:

1. Create a work IODF from the current 2827 production IODF.
2. Repeat the 2827 processor that is being replaced.
3. Observe the CF link messages for later reference.
4. Delete unsupported items for the repeated 2827.
5. Change the repeated 2827 to a 2964, and delete the 2827.
6. Redefine all CF connections to other processors and any Internal CF required connections.
7. Define any additional CHPIDs, control units, and devices you want to add during the upgrade.
8. Overdefine channel paths where needed.
9. Save and restore the OSA configuration information.
10. Build a 2964 validated work IODF.
11. Create an IOCP statements file and file transfer to your workstation. This step can be performed with HCM.
12. Import CFReport and IOCP statements into the CMT.
13. Perform hardware resolution and PCHID/CHPID availability.

14. Create configuration reports for yourself and the IBM System Service Representative.
15. Import IOCP statements that are updated with PCHIDs back into the validated work IODF.
16. Build the production IODF.
17. Remotely write the IOCP to an IOCDS on the 2964 Support Element. If this action is not possible, copy the IOCP statements to a USB memory flash drive and run the Stand-Alone Input/Output Configuration Program to load the IOCP statements to an IOCDS on the 2964 Support Element.
18. Build Reset, Image, and Load Profiles if required.
19. Perform a power-on reset (activate) of the 2964.

5.2.1 Creating the work IODF from the current 2827 production IODF

We use HCD to make a work IODF, but first we start from the current production IODF that contains the 2827 processor you are upgrading (for example, `SYS6.IODFB3`).

5.2.2 Repeating the 2827 processor to be replaced

To repeat the 2827 processor in HCD, complete the following steps:

1. From the main HCD panel, select option **1.3. Processor List**.
2. In the Processor List (Figure 5-2), enter **r** (for repeat) next to the 2827 that you want to upgrade and press Enter.

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
r SCZP401  2827     H43      LPAR  00B8D72827 Helix
```

Figure 5-2 HCD - 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 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.IODFB5.WORK`).

5. The Repeat Processor panel opens (Figure 5-3). Enter the Processor ID of the new 2964 (in this example, SCZP501), keep all the other fields unchanged, and press Enter.

```
+---------------------------- Repeat Processor -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . . . . . . SCZP501_                                   |
|                                                                           |
| Processor type . . . . . . . : 2827                                       |
| Processor model  . . . . . . : H43                                        |
| Configuration mode . . . . . : LPAR                                       |
|                                                                           |
| Serial number  . . . . . . . . 00B8D72827                                 |
| Description  . . . . . . . . . Helix                                      |
|                                                                           |
| Specify SNA address only if part of an System z cluster:                  |
|                                                                           |
| Network name . . . . . . . . . USIBMSC   +                                |
| CPC name . . . . . . . . . . . SCZP401   +                                |
|                                                                           |
| Local system name  . . . . . . SCZP401                                    |
+-----------------------------------+                                       |
| New IODF SYS6.IODFB5.WORK defined. | F4=Prompt   F5=Reset    F9=Swap      |
+-----------------------------------+                                       |
```

Figure 5-3 HCD - Repeat Processor: defining new Processor ID

5.2.3 CF Link information messages

You might receive severity messages (E, I, or W). As shown in Figure 5-4, `CBDG441I`, severity I messages are displayed in the example because the CF Link CHPIDs were not copied to the 2964 definition.

```
+------------------------------- Message List --------------------------------+
|   Save  Query  Help                                                        |
| -------------------------------------------------------------------------- |
|                                                            Row 232 of 242  |
| Command ===> ___________________________________________ Scroll ===> CSR   |
|                                                                            |
| Messages are sorted by severity. Select one or more, then press Enter.     |
|                                                                            |
| / Sev Msg. ID  Message Text                                                |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              3.D5 of processor SCZP401 and channel path 0.D7 of          |
| #              processor SCZP401 is not copied.                            |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              3.D6 of processor SCZP401 and channel path 2.D4 of          |
| #              processor SCZP401 is not copied.                            |
| _ I   CBDG441I The coupling facility connection between channel path       |
| #              3.D7 of processor SCZP401 and channel path 2.D5 of          |
| #              processor SCZP401 is not copied.                            |
| _ I   CBDG271I Requested action on object SCZP401 successfully             |
| #              processed.                                                  |
| **************************** Bottom of data ****************************** |
+----------------------------------------------------------------------------+
```

Figure 5-4 HCD - Message List: showing CBDG441I

To resolve this problem, complete the following steps:

1. Scroll until you reach the end of the messages and see the `CBDG271I requested action on object SCZP401 successfully processed` message.
2. Press PF3 or PF12 to continue. As shown in Figure 5-5, there is an extra 2827 processor named SCZP501.

```
 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2827     H43      LPAR  00B8D72827 Helix
```

Figure 5-5 HCD - Processor List: repeated processor

5.2.4 Changing the 2827 to a 2964 and deleting the 2827

You can either keep the original copy of the 2827 (SCZP401) or delete it from the IODF. In this example, keep it in the IODF for a few more steps.

To change the 2827 to a 2964, complete the following steps:

1. Enter `c` (for change) next to SCZP501 to change the 2827 to a 2964 and press Enter. The Change Process Definition panel opens (Figure 5-6).
2. Make the following updates and press Enter:
   - Update `Processor type` to `2964`.
   - Update `Processor model` to `N63`.
   - Update the 2827 part of the serial number to 2964 (that is, `0B8D72827` to `0B8D72964`).
   - Update the Description if you want

```
+----------------------- Change Processor Definition ------------------------+
|                                                                            |
|                                                                            |
| Specify or revise the following values.                                    |
|                                                                            |
| Processor ID . . . . . . . . : SCZP501                                     |
| Support level:                                                             |
| XMP, 2827 GA2 support                                                      |
| Processor type . . . . . . . . 2964     +                                  |
| Processor model  . . . . . . . N63      +                                  |
| Configuration mode . . . . . . LPAR     +                                  |
|                                                                            |
| Serial number  . . . . . . . . 00B8D72964  +                               |
| Description  . . . . . . . . . Sphinx                                      |
|                                                                            |
| Specify SNA address only if part of an System z cluster:                   |
|                                                                            |
| Network name . . . . . . . . . USIBMSC  +                                  |
| CPC name . . . . . . . . . . . SCZP401  +                                  |
|                                                                            |
| Local system name  . . . . . . SCZP401                                     |
|  F1=Help     F2=Split    F3=Exit     F4=Prompt   F5=Reset    F9=Swap       |
| F12=Cancel                                                                 |
+----------------------------------------------------------------------------+
```

Figure 5-6 HCD - Change Processor Definition: changing repeated processor

3. The Update Channel Path Identifiers panel opens (Figure 5-7). No changes are made in this example.

```
+-------------------- Update Channel Path Identifiers ---------------------+
|                                                              Row 1 of 88 |
| Command ===> __________________________________________ Scroll ===> CSR   |
|                                                                          |
| Specify any changes to the channel path identifiers in the list below.   |
|                                                                          |
| Processor ID . . . . : SCZP501      Helix                                |
| Channel Subsystem ID : 0                                                 |
|                                                                          |
| CHPID  Type  Side  Until CHPID   New CHPID +                             |
| 00     OSD         __            00                                      |
| 01     OSC         __            01                                      |
| 02     OSD         __            02                                      |
| 03     OSD         __            03                                      |
| 04     OSD         __            04                                      |
| 06     OSD         __            06                                      |
| 07     OSD         __            07                                      |
| 0C     OSD         __            0C                                      |
| 0D     OSC         __            0D                                      |
| 12     OSD         __            12                                      |
| 13     OSD         __            13                                      |
|  F1=Help       F2=Split      F3=Exit       F4=Prompt     F5=Reset        |
|  F7=Backward   F8=Forward    F9=Swap      F12=Cancel                     |
+--------------------------------------------------------------------------+
```

Figure 5-7 HCD - Update Channel Path Identifiers: not changed

4. Press Enter *for each Channel Subsystem ID*.

   HCD displays the message shown in Figure 5-8 because the 2964 has more configuration resources than the 2827 processor definition.

```
+-----------------------------------------------------------------+
| Definition of processor SCZP501 has been extended to its maximum |
| configuration.                                                  |
+-----------------------------------------------------------------+
```

Figure 5-8 HCD - Processor List: maximum configuration extended

5. Press Enter. The repeated 2827 processor is successfully changed to a 2964-N63 as shown in Figure 5-9.

```
 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  00B8D72964 Sphinx
```

Figure 5-9 HCD - Processor List: changed processor

5.2.5 Deleting the 2827 processor definition

Now that the 2827 was repeated and changed to a 2964, the original 2827 definition (SCZP401) must be deleted so that the required CF Links can be restored.

To delete the 2827 processor definition, complete the following steps:

1. Enter d (for delete) next to the SCZP401 processor in the Processor List (Figure 5-10).

```
 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
d SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  00B8D72964 Sphinx
```

Figure 5-10 HCD - Processor List: deleting processor

2. Press Enter to confirm the deletion of the processor (Figure 5-11).

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP501  2964     N63      LPAR  00B8D72964 Sphinx
```

Figure 5-11 HCD - Processor List: processor deleted

5.2.6 Reconnecting the CF channel paths that were not migrated

Manually redefine the CF Links that you want from the SCZP501 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 2827. Alternatively, you can make a note of all CBDG441I error messages that you received in 5.2.5, “Deleting the 2827 processor definition” on page 175.

5.2.7 Defining more I/O

Define any additional CHPIDs, control units and devices, FICON channel-to-channel connections (FCTCs), and so on, that you might be adding into the 2964 during the replace. This section shows use of HCD and HCM to define the following CHPIDs and PCIe functions:

- Defining a CF link over CHPID type CS5 using HCD
- Defining a CF link over CHPID type CS5 using HCM
- Defining a RoCE PCIe function using HCD
- Defining a RoCE PCIe function using HCM
- Defining a zEDC EXPRESS PCIe function by using HCD
- Defining a zEDC EXPRESS PCIe function by using HCM

Defining a CF link over CHPID type CS5 using HCD

A new CHPID type (CS5) is introduced with the z13 and uses a new PCIe-O SR ICA hardware card for the z13 named `PCIe-O SR Integrated Coupling Adapter FC 0172`.

It runs at 8 GBps and up to a length of 150 meters. Presently, its connectivity is supported only between a z13 and another z13.

The HCD definition process is similar to CHPIDs of type=CIB except that the type=CS5.

In this example, we define these items:

- CHPID=0.86 to AID/Port ID=27/1 to processor=SCZP501
- CHPID=0.96 to AID/Port ID=27/2 to processor=SCZP502
- A CF link between these two CS5 CHPIDs

Complete these steps:

1. From the main HCD panel, select option **1.3. Processor List**. Enter `s` (work with attached channel paths) next to the first CPC (SCZP501) to which you want to define the CS5 CHPID and press Enter.
2. From the Channel Subsystem List, enter `s` (work with attached channel paths) next to the CSSID you want. In our example, we use CSSID=0. Press Enter.
3. Enter **`add`** on the command line in the Channel Path List panel to add a new channel and press Enter.
4. Enter the following values in the Add Channel Path panel (Figure 5-12) and press Enter.
   - Channel path ID: `86`
   - Channel path type: `CS5`
   - Operational mode: `SHR`

```
--------------------------- Add Channel Path --------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP501           Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 86    +           Channel ID ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +
Description  . . . . . . CS5 CHPID definition___________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 5-12 HCD - Add Channel Path: defining a CS5 CHPID

5. Enter these values in the Specify HCA Attributes panel (Figure 5-13) and press Enter:
   – Adapter ID of the HCA: 27
   – Port on the HCA: 1

```
----------------- Specify HCA Attributes -------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 1   +
```

Figure 5-13 HCD - Specify HCA Attributes: defining CS5 HCA attributes

6. Select the required Access LPARs for CHPID access list. In our example, we use LPAR A0C (CF) and A01, A02 and A03 (OS). Then press Enter. See Figure 5-14.

```
-------------------------- Define Access List ---------------------------
                                                              Row 1 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 86     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      OS
_ 0      A0B              B      CF/OS CHPID holder
/ 0      A0C              C      CF    CF7A
_ 0      A0D              D      CF    CF7B
_ 0      A0E              E      CF    CF7C
_ 0      A0F              F      CF    CF7D
/ 0      A01              1      OS    SC76
/ 0      A02              2      OS    SC74
/ 0      A03              3      OS    SC75
_ 0      A04              4      OS    ZZOVM1
```

Figure 5-14 HCD - Define Access List: selecting partition for CS5 CHPID access

7. Select any Candidate LPARs for the CHPID access list. In our example, we do not select any candidate LPARs. Press Enter.
8. From the main HCD panel, select option **1.3. Processor List**. Enter s (work with attached channel paths) next to the second CPC (SCZP502) to which you want to define the CS5 CHPID and press Enter.
9. From the Channel Subsystem List, enter s (work with attached channel paths) next to the CSSID. In our example, we use CSSID=0. Press Enter.
10. Type **add** on the command line in the Channel Path List panel to add a new channel and press Enter.

11. Enter these values in the Add Channel Path panel (Figure 5-15) and press Enter.

– Channel path ID: 96
– Channel path type: CS5
– Operational mode: SHR

```
--------------------------- Add Channel Path ------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP502      test machine definition
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 96    +           Channel ID  ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +
Description  . . . . . . CS5 CHPID definition____________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 5-15 HCD - Add Channel Path: defining a CS5 CHPID

12. Enter these values in the Specify HCA Attributes panel (Figure 5-16) and press Enter.

– Adapter ID of the HCA: 27
– Port on the HCA: 2

```
---------------- Specify HCA Attributes ------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 2   +
```

Figure 5-16 HCD - Specify HCA Attributes: defining CS5 HCA attributes

13. Select the required Access LPAR for CHPID access list. In our example, we use LPAR A0D (CF) and A01, A02, and A03 (OS). Then press Enter. See Figure 5-17.

```
-------------------------- Define Access List ---------------------------
                                                              Row 1 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 96     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      CF/OS
_ 0      A0B              B      CF/OS
_ 0      A0C              C      CF
/ 0      A0D              D      CF
_ 0      A0E              E      CF/OS
_ 0      A0F              F      CF/OS
/ 0      A01              1      OS
/ 0      A02              2      OS
/ 0      A03              3      OS
_ 0      A04              4      CF/OS
```

Figure 5-17 HCD - Define Access List: selecting partition for CS5 CHPID access

14. Select any Candidate LPARs for the CHPID access list. In our example, we do not select any candidate LPARs. Press Enter.

15. Return to the first CHPID definition (0.86 on processor SCZP501) and type `f` (connect CF channel paths) next to CHPID 0.86 and press Enter. See Figure 5-18.

```
                              Channel Path List    Row 91 of 117 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f 86    27/1  CS5   SHR   __  __ __   N   No  CS5 CHPID definition
_ 87    34/1  CS5   SHR   __  __ __   N   No  P Fram
_ 90    08/2  CIB   SHR   __  __ __   Y   No  12x P201 T/O
_ 94    0B/2  CIB   SHR   __  __ __   Y   No  1x
_ 95    0F/2  CIB   SHR   __  __ __   Y   No  1x
```

Figure 5-18 HCD - Channel Path List: connecting a CF channel path

16. Type p (CF channel path connectivity list) next to CHPID 0.86 and press Enter. See Figure 5-19.

```
                     CF Channel Path Connectivity List            Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N     SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N     SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N     SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N     SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N     SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N     SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
p 86  27/1  Y  CS5  SHR  N
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N     SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 5-19 HCD - CF Channel Path Connectivity List: creating a CF link

17. Enter the following values in the Connect to CF Channel Path panel (Figure 5-20) and press Enter.

– Destination processor ID: SCZP502
– Destination channel subsystem ID: 0
– Destination channel path ID: 96

```
--------------- Connect to CF Channel Path --------


Specify the following values.

Source processor ID  . . . . . : SCZP501
Source channel subsystem ID  . : 0
Source channel path ID . . . . : 86
Source channel path type . . . : CS5

Destination processor ID . . . . . . SCZP502_  +
Destination channel subsystem ID . . 0   +
Destination channel path ID  . . . . 96  +

Timing-only link . . . . . . . . . . No
```

Figure 5-20 HCD - Connect to CF Channel Path: specifying the destination attributes

18. Accept or override the Control Unit and Device numbers for processor SCZP501, then press Enter. See Figure 5-21.

```
---------------- 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 . . . . . . . : SCZP501
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 86         Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF6  +

Device number  . . . . . . . FE3D
Number of devices  . . . . . 8
```

Figure 5-21 HCD - Add CF Control Unit and Devices: accept or override default values

19. Accept or override the Control Unit and Device numbers for processor SCZP502, then press Enter. See Figure 5-22.

```
---------------- 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 . . . . . . . : SCZP502
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 96         Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF5  +

Device number  . . . . . . . FE2D
Number of devices  . . . . . 8
```

Figure 5-22 HCD - Add CF Control Unit and Devices: accept or override default values

HCD returns to the CF Channel Path Connectivity List panel (Figure 5-23) where you can see that the CF link is now defined.

```
                     CF Channel Path Connectivity List             Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------    -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ    Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N      SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N      SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N      SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N      SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N      SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N      SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
_ 86  27/1  Y  CS5  SHR  N      SCZP502.0  96  27/2  Y  CS5  SHR    CFP   8
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N      SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 5-23 HCD - CF Channel Path Connectivity List: CF link now created

Defining a CF link over CHPID type CS5 using HCM

A CHPID type (CS5) was introduced with the IBM z13 and uses a new PCIe-O SR ICA hardware card for the IBM z13 named *FC 0172 PCIe-O SR Integrated Coupling Adapter*.

It runs at 8 GBps and up to a length of 150 meters. Presently, its connectivity is supported only between one IBM z13 and another IBM z13.

The HCM definition process is similar to CHPIDs of type=CIB except that the type=CS5.

In this example, we define the following items:

- CHPID=0.86 to AID/Port ID=27/1 to processor=SCZP501
- CHPID=0.96 to AID/Port ID=27/2 to processor=SCZP502
- A CF link between these two CS5 CHPIDs

Complete these steps:

1. From the HCM main page, click **Edit** → **Processor**.
2. Select the first CPC (SCZP501), and then click **OK**.
3. Click once on the desired CSS (CSSID=0), and then click **CHPIDs**.
4. Click **Create**. HCM displays the Create CHPIDs window.
5. Enter the following values in the Create CHPIDs window, then click **OK**. See Figure 5-24 on page 183.
   - Channel path ID: `86`
   - Channel path type: `CS5`
   - Operational mode: `SHR`

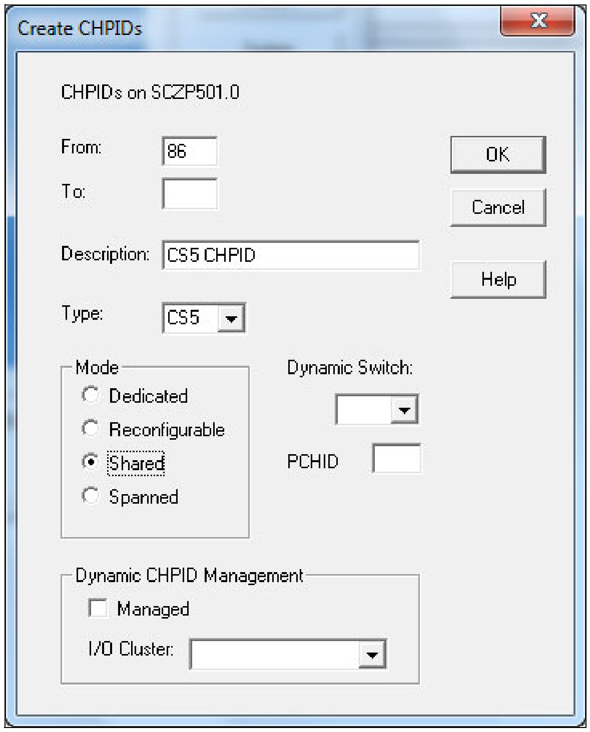

Figure 5-24 HCM - Create CHPIDs: defining a CS5 CHPID

6. Enter the following values in the Host Communication Adapter Attributes window (Figure 5-25), and then click **OK**.
   – Adapter ID of the HCA: `27`
   – Port on the HCA: `1`

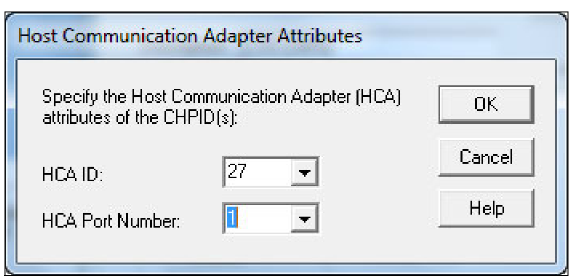

Figure 5-25 HCM - Specify HCA Attributes: defining CS5 HCA attributes

7. Select the required Access and Candidate LPARs for CHPID access list. In our example, we use LPAR A0C (CF) and A01, A02 and A03 (OS), and no Candidate LPARs. Then click **OK**. See Figure 5-26 on page 184.

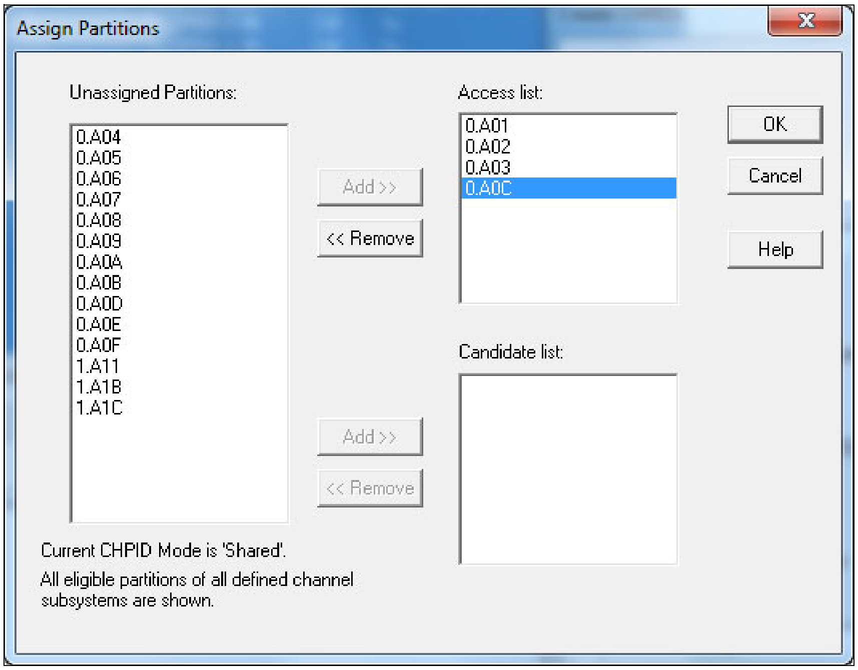

Figure 5-26 HCM - Assign Partitions: selecting partition for CS5 CHPID access

8. Now observe the CS5 CHPID defined in the CHPID summary list (Figure 5-27).

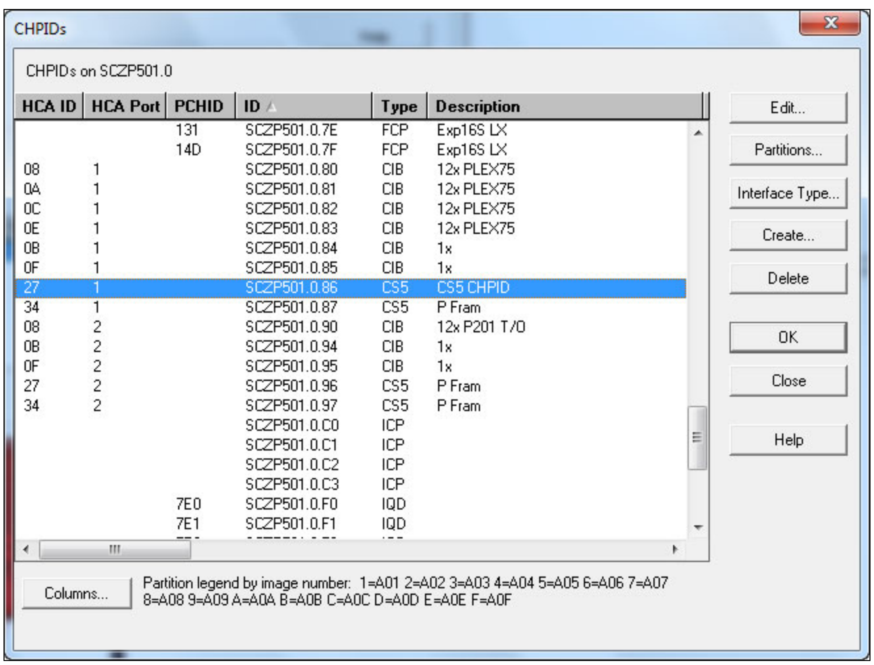

Figure 5-27 HCM - CHPIDs: first CS5 CHPID defined

9. From the HCM main page, click **Edit** → **Processor**.
10. Select the second CPC (SCZP502), and then click **OK**.
11. Click once on the desired CSS (CSSID=0), and then click **CHPIDs**.
12. Click **Create**. HCM displays the Create CHPIDs window.

13. Enter the following values in the Create CHPIDs window, and then click **OK**:
   – Channel path ID: `96`
   – Channel path type: `CS5`
   – Operational mode: `SHR`
14. Enter the following values in the Host Communication Adapter Attributes window, and then click **OK**:
   – Adapter ID of the HCA: `27`
   – Port on the HCA: `2`
15. Select the required Access and Candidate LPARs for CHPID access list. In our example, we use LPAR A0D (CF) and A01, A02 and A03 (OS), and no Candidate LPARs. Click **OK**.
16. Now observe the CS5 CHPID defined in the CHPID summary list (Figure 5-28).

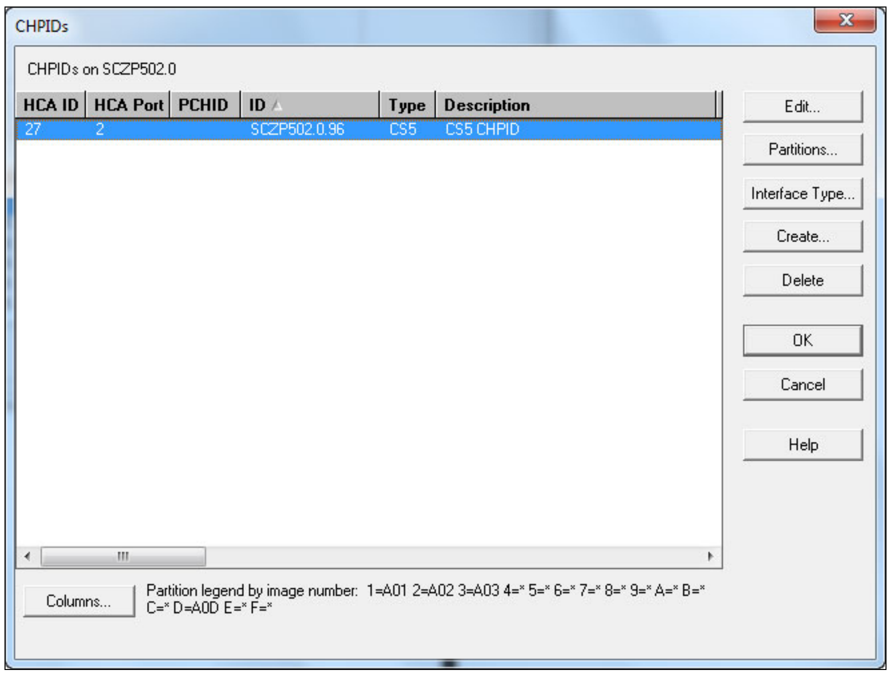

Figure 5-28 HCM - CHPIDs: second CS5 CHPID defined

17. From the HCM main page, click **Create** → **CF connection**.
18. HCM displays the Create Coupling Facility Connection window (Figure 5-29). From the pull-down, select the first CS5 CHPID that we defined (SCZP501 - 0.86).

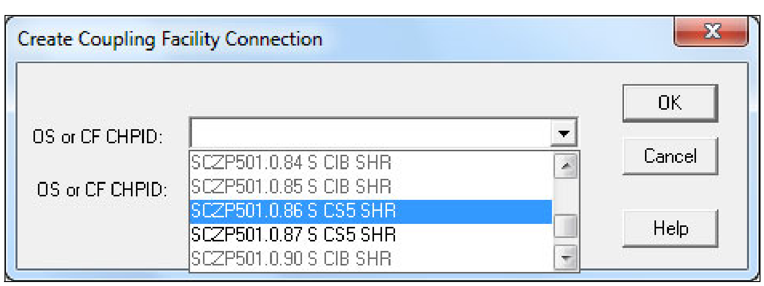

Figure 5-29 HCM - Create Coupling Facility Connection: selecting first CS5 CHPID

19. From the pull-down, select the second CS5 CHPID that we defined (SCZP502 - 0.96), and then click **OK**. See Figure 5-30.

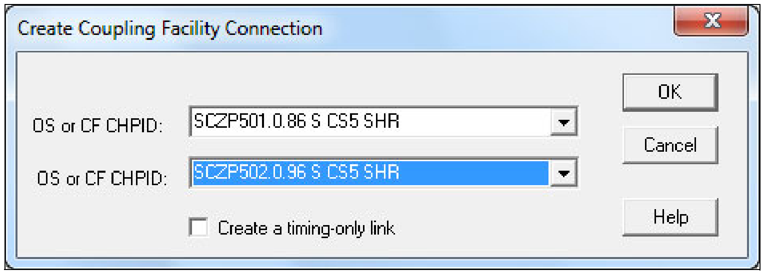

Figure 5-30 HCM - Create Coupling Facility Connection: selecting second CS5 CHPID

20. HCM now suggests control unit and device addresses. Accept or override the control unit and device numbers, then click **OK** (Figure 5-31).

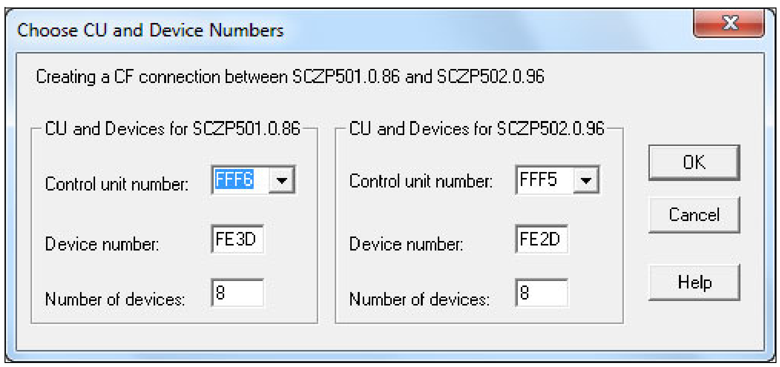

Figure 5-31 HCM - Choose CU and Device Numbers: creating a CF link

21. The CF link is now created. To view it, select the **Locate** → **CF connection** HCM menu option or click one of the CPC CHPID definitions and press F4. See Figure 5-32.

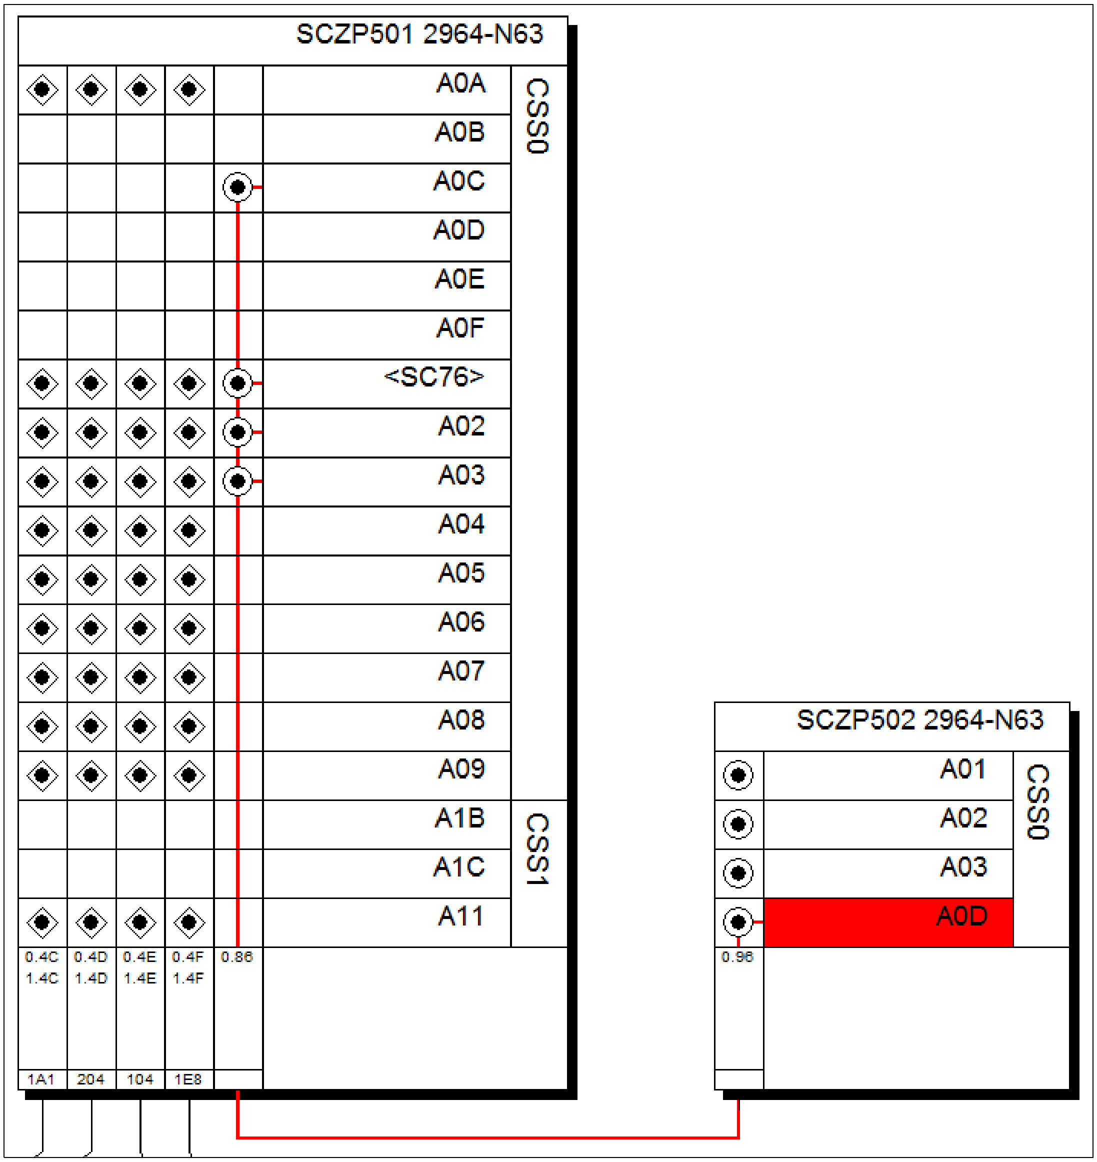

Figure 5-32 HCM - Viewing the CS5 CF link graphically

Defining a RoCE PCIe function using HCD

RDMA over Converged Ethernet (RoCE) uses a new PCIe hardware card for the zEC12 and IBM z13 and is named *10GbE RoCE Express SR* (FC 0411).

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 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 RoCE PCHIDs.

These steps describe only the definition process. However, if you want deeper understanding of how RoCe works, see Chapter 14, "Configure Open System Adaptor (OSA) and RoCE Express" on page 653.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 5-1).

Example 5-1 PCHID Report - RoCE information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                  Comment
19/02/J01       Z22B  20    0411   140                        RG1      NEW
19/12/J01       Z08B  03    0411   208                        RG2      NEW
```

In this example, we define these items:

- PCHID=140 to Function IDs 000 (VF=1) and 001 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1
- PCHID=208 to Function IDs 010 (VF=1) and 011 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1

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 (SCZP502) to which you want to define the RoCE functions and press Enter. See Figure 5-33.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
******************************* Bottom of data ********************************
```

Figure 5-33 HCD - Processor List: adding PCIe functions to a processor

2. To add a new PCIe function, type `add` on the command line in the PCIe Function List panel (Figure 5-34).

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 5-34 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 5-35) and press Enter.
   – Function ID: `000`
   – Type: `ROCE`
   – PCHID: `140`
   – Virtual Function ID: `1`
   – Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID  . . . . : SCZP502      test machine definition

Function ID . . . . . . 000
Type  . . . . . . . . . ROCE         +

PCHID . . . . . . . . . 140
Virtual Function ID . . 1_  +

Description . . . . . . RoCE F=000 VF=1 PCHID=140_______
```

Figure 5-35 HCD - PCIe Function List: adding PCIe functions to a processor

4. Enter the following value in the Add/Modify Physical Network IDs panel (Figure 5-36) and press Enter.
   – Physical network ID 1: `ITSOPNET1`

```
-------------------- Add/Modify Physical Network IDs ------------------


If the PCHID is associated to one or more physical networks, specify
each physical network ID corresponding to each applicable physical port

Physical network ID 1  . . ITSOPNET1_______
Physical network ID 2  . . ________________
Physical network ID 3  . . ________________
Physical network ID 4  . . ________________
```

Figure 5-36 HCD - Add/Modify Physical Network IDs: adding network ID

5. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Press Enter. See Figure 5-37.

```
-------------------------- Define Access List ----------------------------
                                                              Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 5-37 HCD - Define Access List: selecting partition for Function access

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 where you can see the Function now defined. See Figure 5-38.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
******************************* Bottom of data ********************************
```

Figure 5-38 HCD - PCIe Function List: Function now created

7. Now define the other Function IDs according to the example so far. See Figure 5-39.

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

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
_ 001    140    2    ROCE            RoCE F=000 VF=2 PCHID=140
_ 010    208    1    ROCE            RoCE F=000 VF=1 PCHID=208
_ 011    208    2    ROCE            RoCE F=000 VF=2 PCHID=208
******************************* Bottom of data ********************************
```

Figure 5-39   HCD - PCIe Function List: All Functions now created

Defining a RoCE PCIe function using HCM

RDMA over Converged Ethernet (RoCE) uses a new PCIe hardware card for the zEC12 and IBM z13 called *10GbE RoCE Express SR* (FC 0411).

To use RoCE hardware functions, Function IDs, Virtual Function IDs, and Physical Network IDs must be defined in HCD to the PCHID that has been assigned to the RoCE hardware cards 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 RoCE PCHIDs.

These steps describe only the definition process. However, if you want deeper understanding of how RoCe works, see Chapter 14, "Configure Open System Adaptor (OSA) and RoCE Express" on page 653.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 5-2).

Example 5-2   PCHID Report - RoCE information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                 Comment
19/02/J01       Z22B  20    0411   140                          RG1    NEW
19/12/J01       Z08B  03    0411   208                          RG2    NEW
```

In this example we define these items:

- PCHID=140 to Function IDs 000 (VF=1) and 001 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1
- PCHID=208 to Function IDs 010 (VF=1) and 011 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1

Complete the following steps:

1. From the HCM main page, click **Edit** → **Processor**.
2. Select the first CPC (SCZP502), and then click **OK**.
3. Click **Edit** → **PCIe** (Figure 5-40 on page 192).

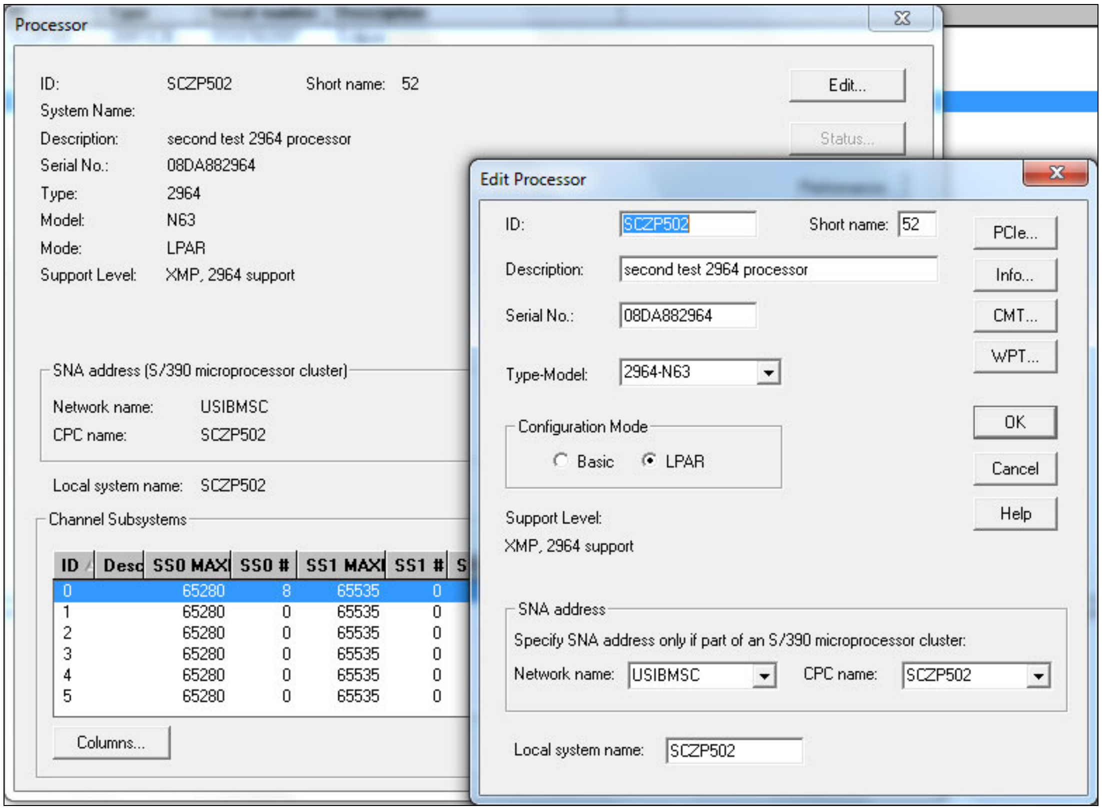

Figure 5-40 HCM - Edit Processor: PCIe functions

4. The PCIe Functions summary window (Figure 5-41) opens. Click **Add**.

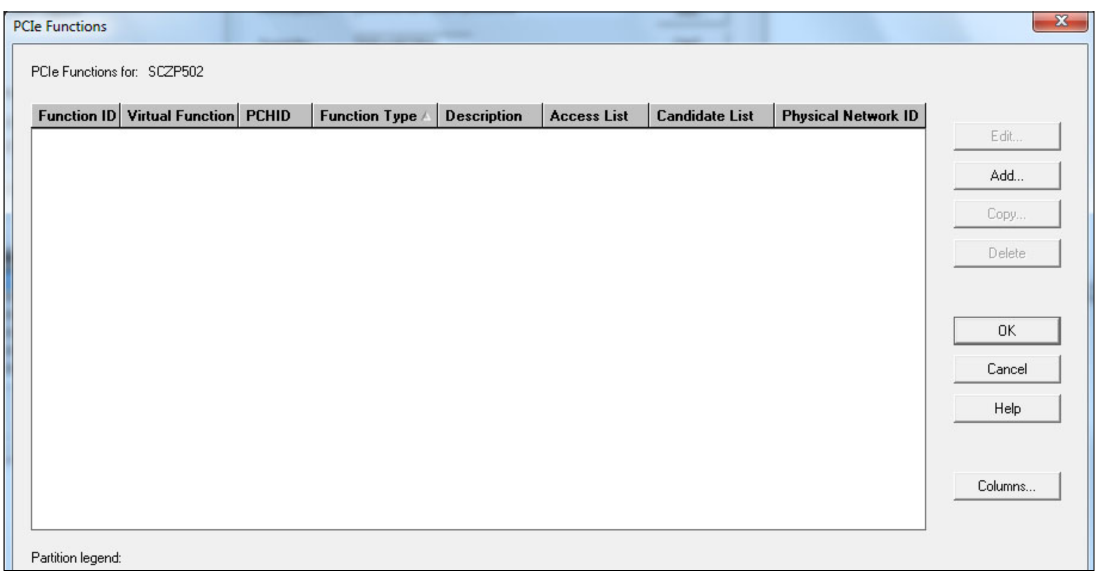

Figure 5-41 HCM - PCIe functions: summary list

5. Enter the following values in the Add PCIe Function window (Figure 5-42 on page 193), and click **OK**:
   – Function ID: `000`
   – Function Type: `ROCE`
   – Virtual Function Number: `1`

– PCHID: `140`
– Description: `optional`
– Physical Network ID 1: `ITSOPNET1`
– LPAR Access List: `0.A01`
– LPAR Candidate List: `0.A02` and `0.A03`

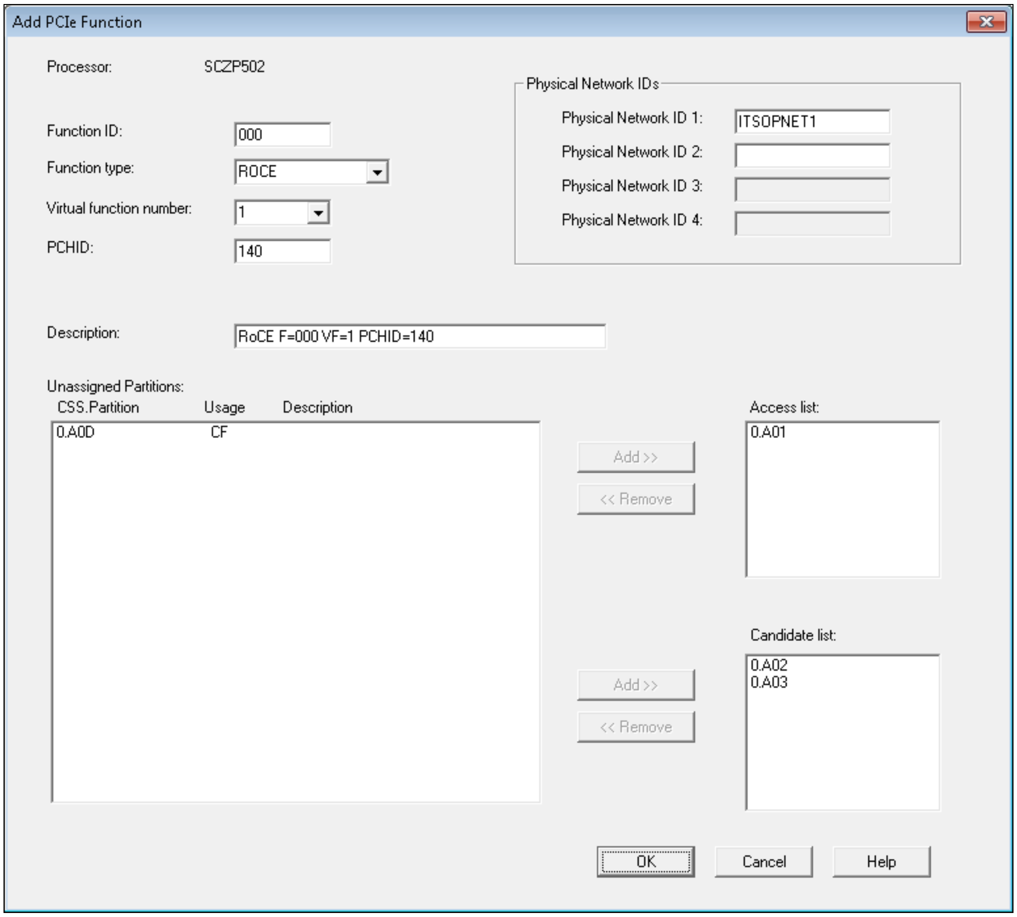

Figure 5-42 HCM - Add PCIe functions: RoCE definition

The RoCE PCIe feature is now listed in the PCIe Function summary list (Figure 5-43).

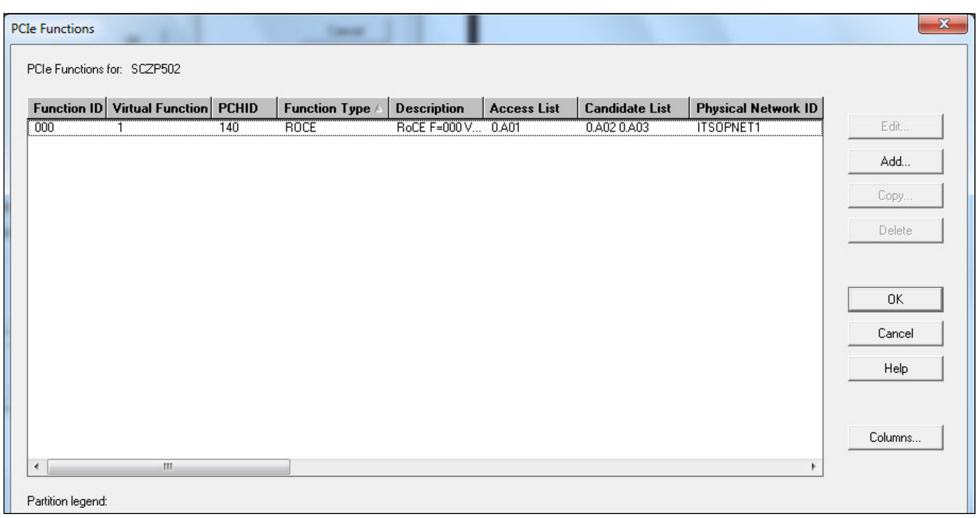

Figure 5-43 HCM - PCIe functions: RoCE PCIe function defined

Defining a zEDC EXPRESS PCIe function by using HCD

The zEnterprise Data Compression (zEDC) uses a new PCIe hardware card for the zEC12 and z13 called *zEDC Express FC 0420.*

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that has been assigned to the zEDC hardware cards 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.

Here we show only the definition process. If you want a deeper understanding of how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 5-3).

Example 5-3 PCHID Report - zEDC information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                Comment
15/14/J01       Z15B  19    0420   1BC                       RG1     NEW
15/05/J01       Z08B  38    0420   27C                       RG2     NEW
```

In this example, we define the following items:

- PCHID=1BC to Function IDs 020, 021 and 022 (VF=1, 2 and 3) on CPC = SCZP502
- PCHID=27C to Function IDs 030, 031 and 032 (VF=1, 2 and 3) on CPC = SCZP502

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 (SCZP502) to which you want to define the zEDC functions and press Enter. See Figure 5-44.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
****************************** Bottom of data ********************************
```

Figure 5-44 HCD - Processor List: adding PCIe functions to a processor

2. Type **add** on the command line in the PCIe Function List panel (Figure 5-45) to add a new PCIe function.

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 5-45 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 5-46) and press Enter.
   - Function ID: `020`
   - Type: `ZEDC-EXPRESS`
   - PCHID: `1BC`
   - Virtual Function ID: `1`
   - Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID  . . . . : SCZP502      test machine definition

Function ID . . . . . . 020
Type  . . . . . . . . . ZEDC-EXPRESS  +

PCHID . . . . . . . . . 1BC
Virtual Function ID . . 1   +

Description . . . . . . zEDC F=020 VF=1 PCHID=1BC_______
```

Figure 5-46 HCD - PCIe Function List: adding PCIe functions to a processor

4. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Then press Enter. See Figure 5-47.

```
-------------------------- Define Access List ----------------------------
                                                              Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 5-47 HCD - Define Access List: selecting partition for Function access

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 (Figure 5-48) where you can see that the Function is now defined.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
******************************* Bottom of data ********************************
```

Figure 5-48 HCD - PCIe Function List: Function now created

6. Now define the other Function IDs according to this example. See Figure 5-49.

```
                              PCIe Function List      Row 1 of 6 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
_ 021    1BC    2    ZEDC-EXPRESS    zEDC F=021 VF=2 PCHID=1BC
_ 022    1BC    3    ZEDC-EXPRESS    zEDC F=022 VF=3 PCHID=1BC
_ 030    27C    1    ZEDC-EXPRESS    zEDC F=030 VF=1 PCHID=27C
_ 031    27C    2    ZEDC-EXPRESS    zEDC F=031 VF=2 PCHID=27C
_ 032    27C    3    ZEDC-EXPRESS    zEDC F=032 VF=3 PCHID=27C
******************************* Bottom of data ********************************
```

Figure 5-49 HCD - PCIe Function List: All Functions now created

Defining a zEDC EXPRESS PCIe function by using HCM

IBM z Enterprise Data Compression (zEDC) uses a new PCIe hardware card for the zEC12 and z13 named *zEDC Express FC 0420*.

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that has been assigned to the zEDC hardware cards 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.

Here we show only the definition process. If you want a deeper understanding of how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 5-4).

Example 5-4 PCHID Report - zEDC information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                  Comment
15/14/J01       Z15B  19    0420   1BC                          RG1    NEW
15/05/J01       Z08B  38    0420   27C                          RG2    NEW
```

In this example, we define the following items:

- PCHID=1BC to Function IDs 020, 021 and 022 (VF=1, 2 and 3) on CPC = SCZP502
- PCHID=27C to Function IDs 030, 031 and 032 (VF=1, 2 and 3) on CPC = SCZP502

Complete the following steps:

1. From the HCM main page, click **Edit** → **Processor**.
2. Select the first CPC (SCZP502), then click **OK**.
3. Click **Edit** → **PCIe**. See Figure 5-50 on page 198.

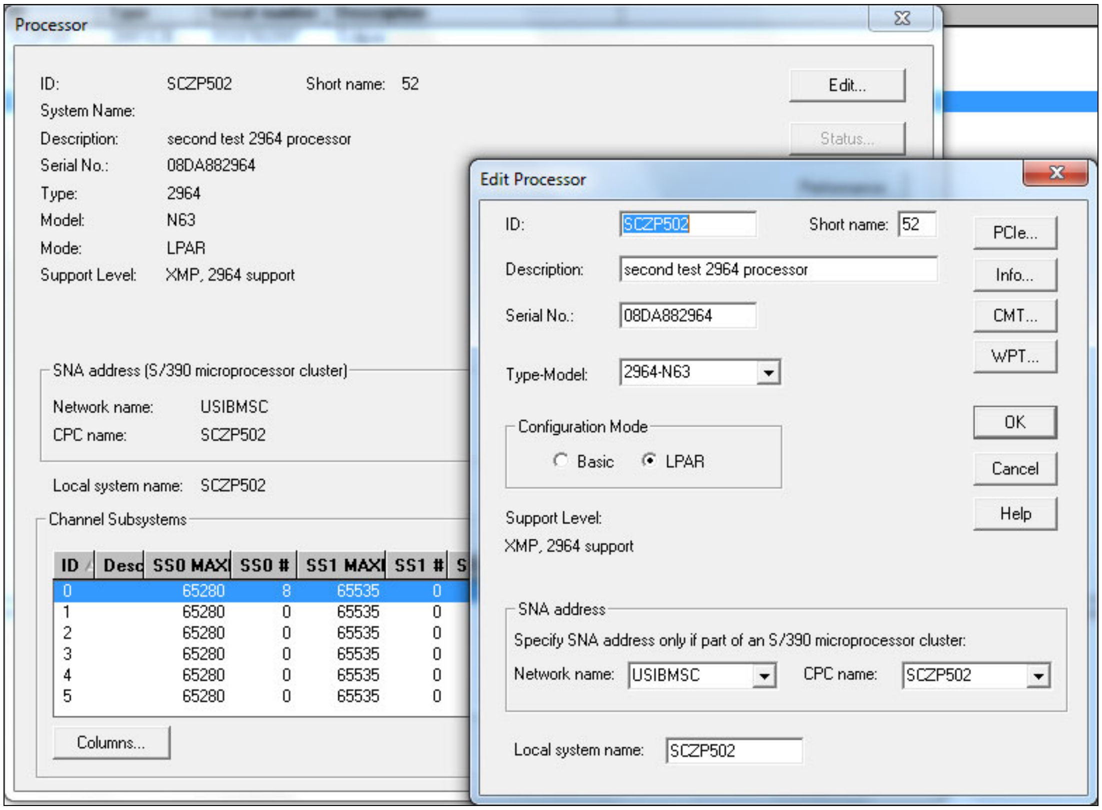

Figure 5-50 HCM - Edit Processor: PCIe functions

4. HCM displays the PCIe Functions summary window (Figure 5-51). Click **Add**.

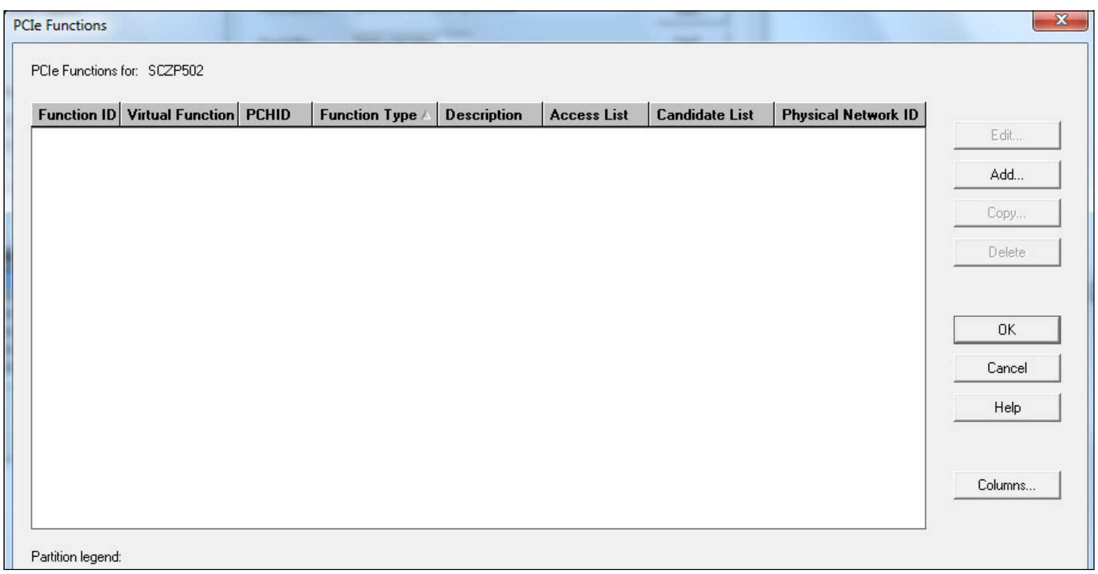

Figure 5-51 HCM - PCIe functions: summary list

5. Enter the following values in the Add PCIe Function window (Figure 5-52 on page 199), and click **OK**.
   - Function ID: `020`
   - Function Type: `ZEDC-EXPRESS`
   - Virtual Function Number: `1`

- PCHID: `1BC`
- Description: `optional`
- LPAR Access List: `0.A01`
- LPAR Candidate List: `0.A02` and `0.A03`

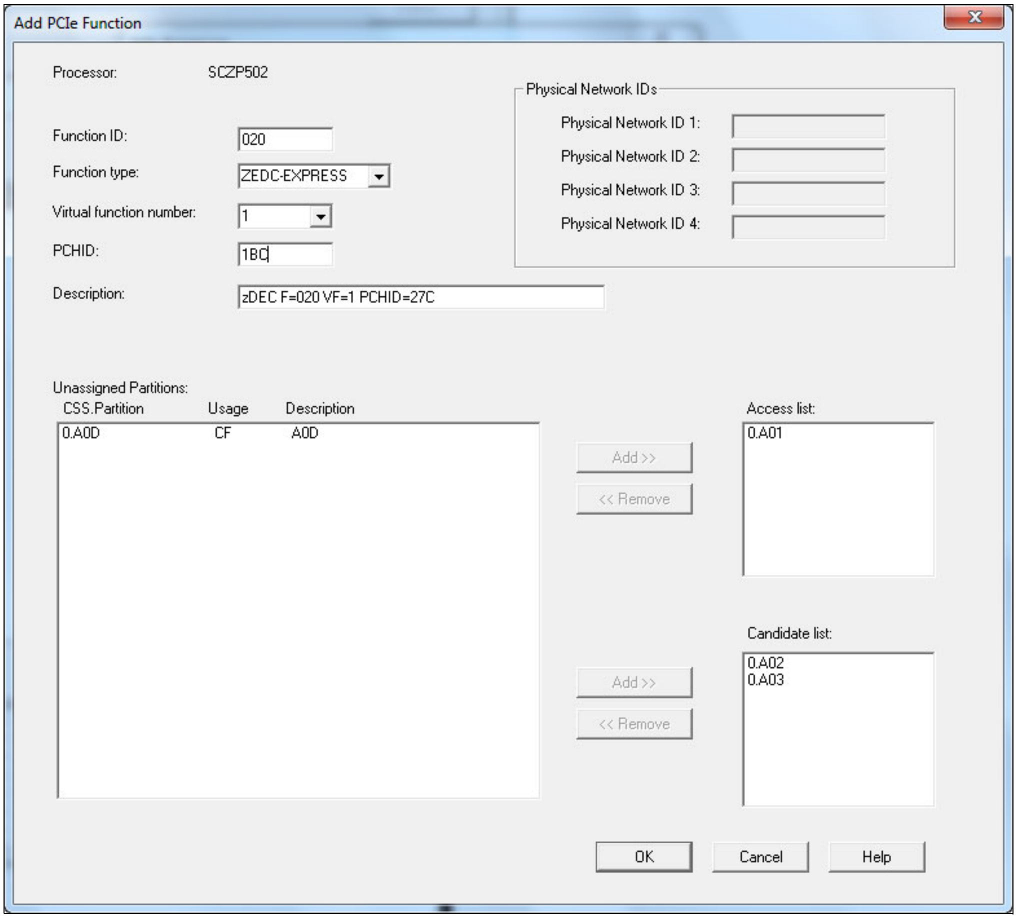

Figure 5-52 HCM - Add PCIe functions: zEDC definition

The zEDC PCIe feature is now listed in the PCIe Function summary list (Figure 5-53).

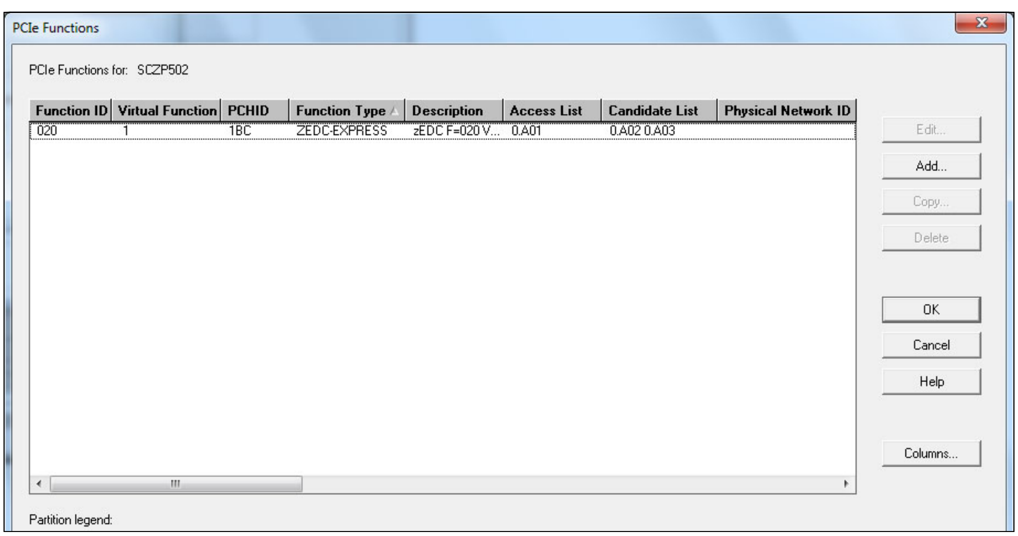

Figure 5-53 HCM - PCIe functions: zEDC PCIe function defined

5.2.8 Overdefining channel paths on an XMP processor

Sometimes you must define a channel path that is not physically installed on the processor. This definition is useful if you are planning to add more channel cards to the processor in the future and want to have the definitions in the IODF before the hardware is installed.

HCD allows you to overdefine CHPIDs by using an asterisk (*) for the PCHID value. An overdefined CHPID must adhere to all validation rules, but it is not taken into account by an IOCDS download. Also, it is not included in the IOCP statements, in a CONFIGxx member, or during dynamic activation.

If a control unit contains only CHPIDs with a PCHID value of an asterisk (*), the whole control unit (including any attached devices) is omitted from the configuration to be activated.

When you install the channel path later, you must edit the CHPID and replace the * with its valid PCHID.

> **Remember:** This is not the case for CIB type CHPIDs, where these CHPIDs have connections to other CIB type CHPIDs. Therefore, HCD allows you to define CIB type CHPIDs as overdefined only if they are unconnected.
>
> Overdefining is also supported for CS5 type CHPID definitions.

The 2964 production IODF can then be activated dynamically and the PCHID, CHPID, and control unit definitions become available to the operating system.

Figure 5-54 shows what the CHPID/PCHID definitions look like before they are defined as overdefined. Press PF20 (right) in the Channel Path List.

```
                              Channel Path List     Row 1 of 107 More: <
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    1A3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    553   FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    153   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 5-54 HCD - Channel Path List: reserving CHPIDs

Figure 5-55 shows how the CHPID/PCHID definitions look after they are defined as overdefined.

```
                              Channel Path List    Row 22 of 109 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                  I/O Cluster   --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    *     FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    *     FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    *     FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 5-55 HCD - Channel Path List: overdefined CHPIDs

5.3 OSA: Saving and restoring configuration data

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

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

See 14.1, “Configure OSE channel with Open System Facility (OSA/SF) on the HMC” on page 654 for details of how to save and restore any Open Systems Adapter (OSA) configuration data such the OSA Address Table (OAT).

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

See 14.5, “Configure OSA-Express Integrated Console Controller (OSA-ICC)” on page 686 if you are unfamiliar with the exporting and importing process for OSA-ICC Server and Session configuration data.

5.3.3 Using OSA Advanced Facilities to set OSA parameters

See 14.6, “Setting OSA parameters by using OSA Advanced Facilities” on page 702 for the process of changing the OSA port speed or MAC addresses.

5.4 HCD: Validating the 2964 work IODF

To validate the 2964 work IODF by using 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 to HCD option **6.4. View I/O Definition File Information**. The IODF type is now indicated as `Work - Validated` (Figure 5-56).

```
+---------------- View I/O Definition File Information -----------------+
|                                                                       |
|                                                                       |
| IODF name  . . . . . . : 'SYS6.IODFB5.WORK'                           |
| IODF type  . . . . . . : Work - Validated                             |
| IODF version . . . . . : 5                                            |
|                                                                       |
| Creation date  . . . . : 2014-11-03                                   |
| Last update  . . . . . : 2014-11-04  17:27                            |
|                                                                       |
| Volume serial number . : IODFPK                                       |
| Allocated space  . . . : 4000    (Number of 4K blocks)                |
| Used space . . . . . . : 2188    (Number of 4K blocks)                |
|    thereof utilized (%)  84                                           |
| Activity logging . . . : No                                           |
| Multi-user access  . . : No                                           |
| Backup IODF name . . . :                                              |
|                                                                       |
| Description  . . . . . :                                              |
|                                                                       |
|                                                                       |
|                                                                       |
|  F1=Help    F2=Split   F3=Exit    F9=Swap   F12=Cancel                |
+-----------------------------------------------------------------------+
```

Figure 5-56 HCD - View I/O Definition File Information: validated work IODF

5.5 Creating the IOCP for the CHPID Mapping Tool

To create the IOCP for the CHPID Mapping Tool, complete the following steps:

> **Tip:** You might prefer to use HCM to create the IOCP statements file and transfer the file to your workstation. You can then start the CHPID Mapping Tool, create an updated IOCP statements file, and transfer the file back to the host.

1. Select HCD option **3. Build IOCP input data set** and press Enter (Figure 5-57).

```
+------ 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   |
```

Figure 5-57 HCD - Activate or Process Configuration Data: building IOCP for SCZP501

2. HCD displays the list of available processors (Figure 5-58). Select the SCZP501 processor by entering a forward slash (/) next to it and pressing Enter.

```
+-------------------------- Available Processors ---------------------------
|                                                               Row 1 of 2
| Command ===> __________________________________________________________
|
| Select one.
|
|   Processor ID  Type     Model    Mode  Description
|   SCZP201       2097     E26      LPAR  Eclipse
| / SCZP501       2964     N63      LPAR  Sphinx
```

Figure 5-58 HCD - 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 (Figure 5-59). Complete the following fields:
   - Title1: `IODFB5`
   - IOCP input data set: `'SYS6.IODFB5.IOCPIN.SCZP501'`
   - 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 . . . . . . . . . : 'SYS6.IODFB5.WORK'                          |
| Processor ID  . . . . . . . : SCZP501                                     |
| Title1 . IODFB5_______________________________________________________    |
| Title2 : SYS6.IODFB5.WORK - 2014-11-04 17:27                              |
|                                                                           |
| IOCP input data set                                                       |
| 'SYS6.IODFB5.IOCPIN.SCZP501'____________________________                  |
| Input to Stand-alone IOCP?  Yes  (Yes or No)                              |
|                                                                           |
| Job statement information                                                 |
| //WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H                                 |
| //*                                                                       |
| //*                                                                       |
| //*                                                                       |
| //*                                                                       |
| //*                                                                       |
|                                                                           |
```

Figure 5-59 HCD - 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 (Figure 5-60). This data set is used as input into the CHPID Mapping Tool.

```
ID    MSG1='IODFB5',                                        *
      MSG2='SYS6.IODFB5.WORK - 2014-11-04 17:27',           *
      SYSTEM=(2964,1),LSYSTEM=SCZP401,                      *
      TOK=('SCZP501',00800001DA8729641727492501143O8F00000000,*
      00000000,'14-11-04','17:27:49','........','........')
RESOURCE PARTITION=((CSS(0),(A0A,A),(A0B,B),(A0C,C),(A0D,D),(A*
      0E,E),(A0F,F),(A01,1),(A02,2),(A03,3),(A04,4),(A05,5),(A*
      06,6),(A07,7),(A08,8),(A09,9)),(CSS(1),(A1A,A),(A1B,B),(*
      A1C,C),(A1D,D),(A1E,E),(A1F,F),(A11,1),(A12,2),(A13,3),(*
      A14,4),(A15,5),(A16,6),(A17,7),(A18,8),(A19,9)),(CSS(2),*
      (A2A,A),(A2B,B),(A2C,C),(A2D,D),(A2E,E),(A2F,F),(A21,1),*
      (A22,2),(A23,3),(A24,4),(A25,5),(A26,6),(A27,7),(A28,8),*
      (A29,9)),(CSS(3),(A3A,A),(A3B,B),(A3C,C),(A3D,D),(A3E,E)*
      ,(A3F,F),(A31,1),(A32,2),(A33,3),(A34,4),(A35,5),(A36,6)*
      ,(A37,7),(A38,8),(A39,9)),(CSS(4),(*,1),(*,2),(*,3),(*,4*
      ),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),*
      (*,E),(*,F)),(CSS(5),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6)*
      ,(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)))
FUNCTION FID=000,PCHID=544,PNETID=ITSOPNET1,PART=((A15),(=)), *
      TYPE=ROCE
FUNCTION FID=010,PCHID=5EC,PNETID=ITSOPNET1,PART=((A03),(=)), *
      TYPE=ROCE
FUNCTION FID=020,VF=1,PCHID=578,PART=((A15),(=)),TYPE=ZEDC
FUNCTION FID=021,VF=2,PCHID=578,PART=((A03),(=)),TYPE=ZEDC
```

Figure 5-60 HCD - IOCP input data set: contents (truncated)

Part of the TOK statement is now replaced with dots (Example 5-5).

Example 5-5 HCD - IOCP file (TOK statement)

```
TOK=('SCZP501',00800001DA87296417274925011430 8F00000000,*
      00000000,'14-11-04','17:27:49','........','........')
```

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 only from production IODFs.

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 `SCZP501in.iocp`.

5.6 CMT: Assigning PCHIDs to CHPIDs

The following steps use the output from the previous set of HCD steps (IOCP) and the 2964 order process (CFReport). Use the CHPID Mapping Tool to assign PCHIDs to each of the CHPIDs for the 2964.

For this process, the CMT must be downloaded. For more information about downloading and installing the CMT, see 3.1.5, “CHPID Mapping Tool (CMT)” on page 44. If CMT is already installed, verify that the latest updates are installed. For more information, see the *CHPID Mapping Tool User’s Guide*, GC28-6947-00.

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 PCHID conflict.
4. Process the hardware resolution.
5. Resolve manually the CIB 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 and transfer it back to the host z/OS image.

Requirement: When you upgrade from a 2817 or a 2827 to a 2964, you must use the CHPID Mapping Tool level that supports the 2964.

5.6.1 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. Right-click the Projects pane and select **New** → **Standard CMT Project** (Figure 5-61).

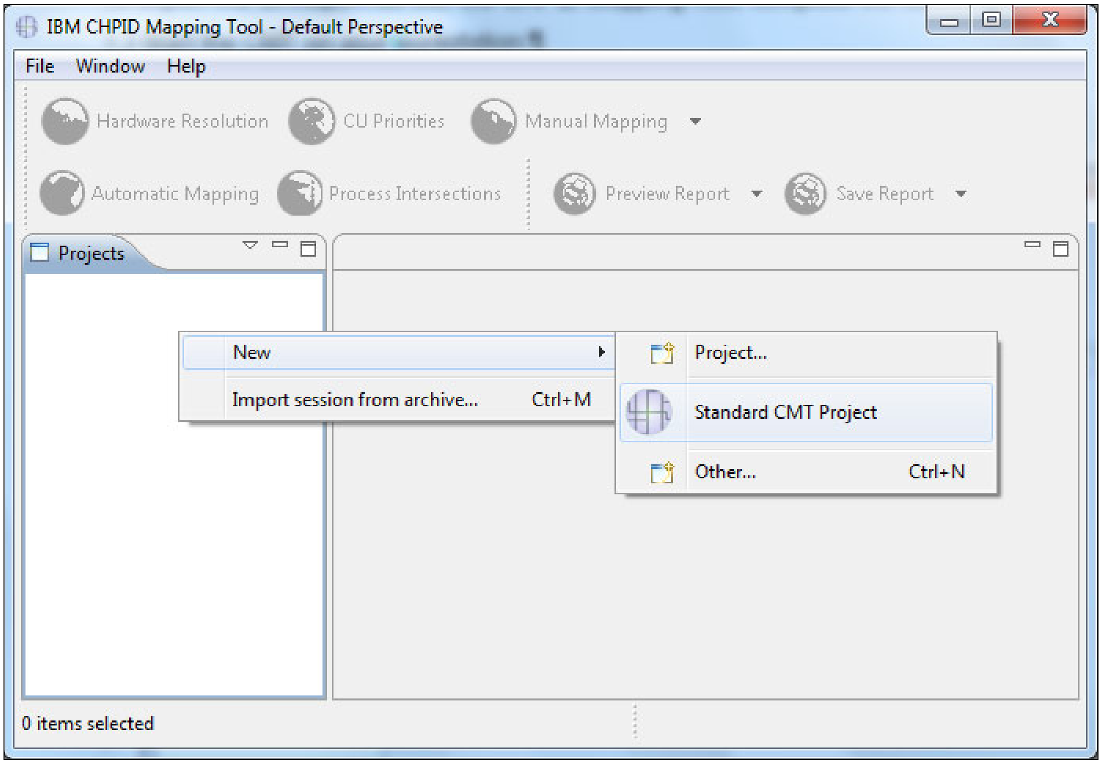

Figure 5-61 CMT - Creating a new CHPID Mapping Tool Project

3. The New CHPID Mapping Tool Project window opens (Figure 5-62 on page 208). Use this window to specify a project name and then click **Next**.

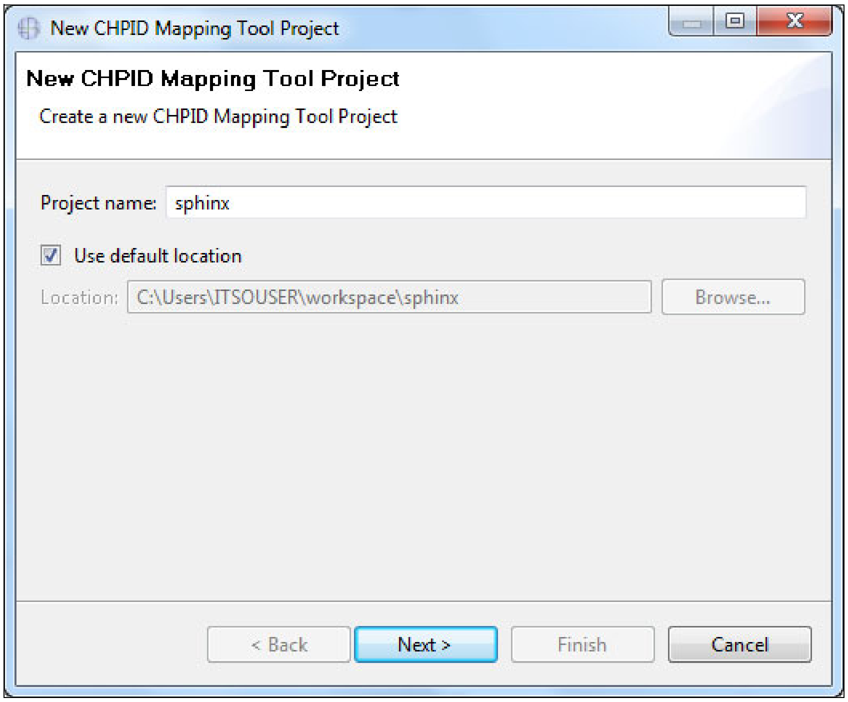

Figure 5-62 CMT - New CHPID Mapping Tool Project name

4. Import the `CFReport` file into the CHPID Mapping Tool by specifying the name in the CFReport file field, and then click **Finish** (Figure 5-63).

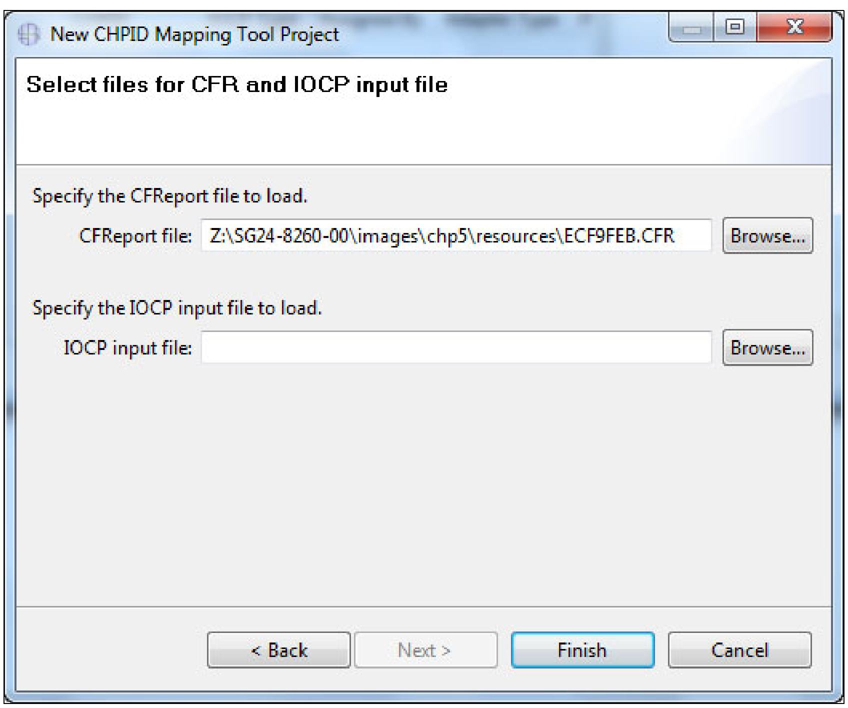

Figure 5-63 CMT - Specifying the CFReport file

If you click **Finish** but did not select an IOCP file, you receive the message shown in Figure 5-64.

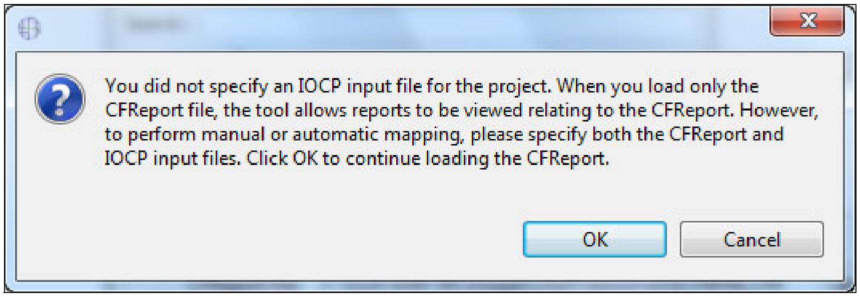

Figure 5-64 CMT - Warning message for not specifying an IOCP file

A window shows the progress of reading the CFReport file (Figure 5-65).

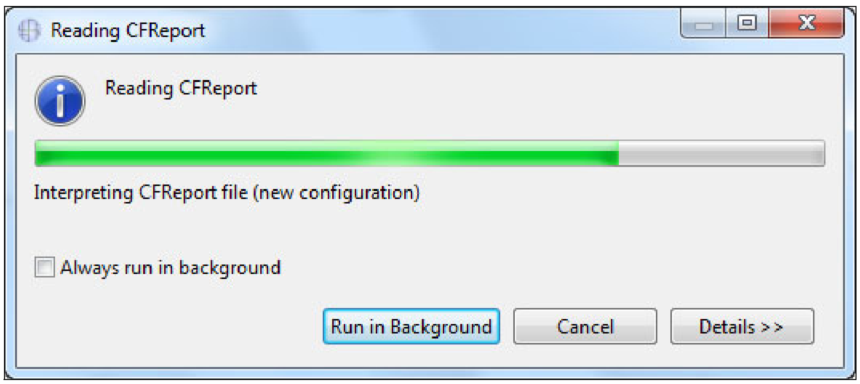

Figure 5-65 CMT - Reading the CFReport file

The information from the CFReport file is displayed in the Hardware pane (Figure 5-66).

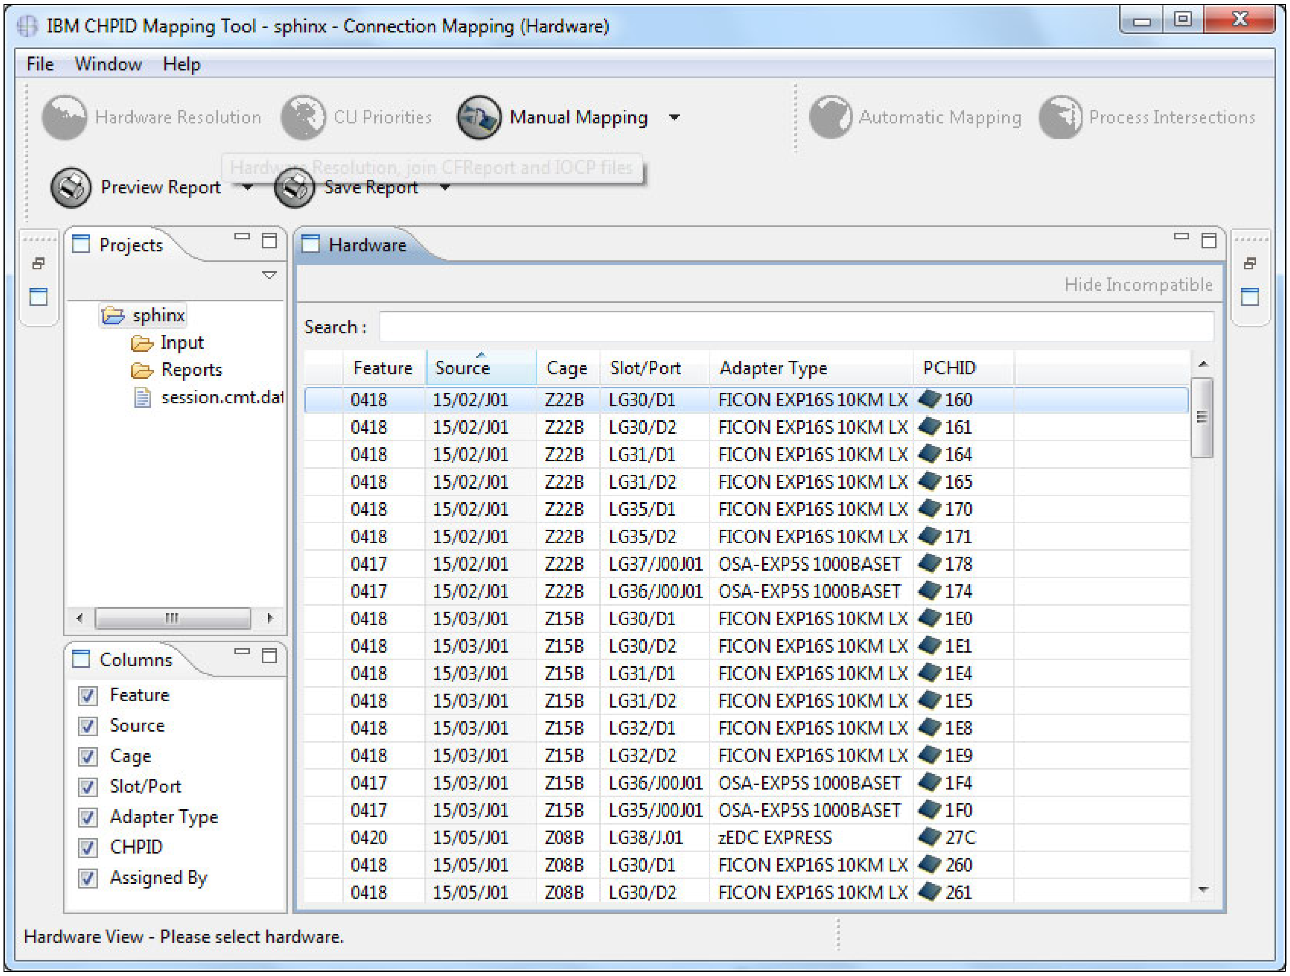

Figure 5-66   CMT - Imported CFReport file

5.6.2 Importing the 2964 IOCP file into the CHPID Mapping Tool

To import the 2964 IOCP file into the CHPID Mapping Tool, complete the following steps:

1. Right-click in the Projects pane and select **Import IOCP input file** (Figure 5-67).

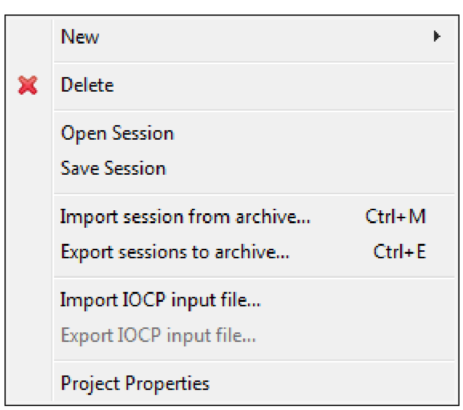

Figure 5-67   CMT - Importing the IOCP file

2. Select the IOCP file on your workstation to import into the CHPID Mapping Tool and click **Finish** (Figure 5-68).

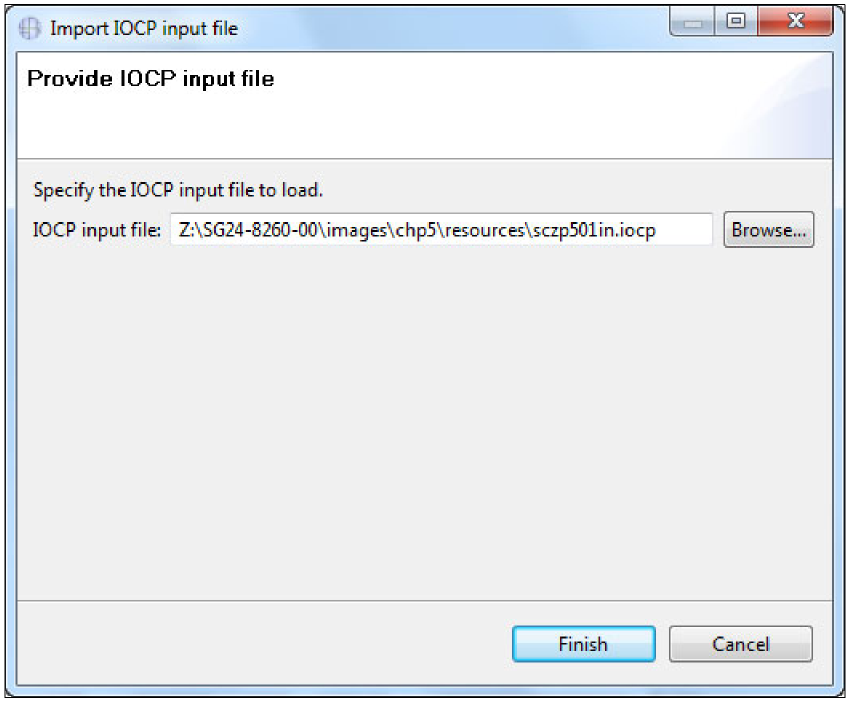

Figure 5-68 CMT - Specifying the IOCP file for import

3. In the Projects pane, under the Input tab, expand the IOCP tab, right-click on the IOCP file, and select **Read Selected IOCP** (Figure 5-69).

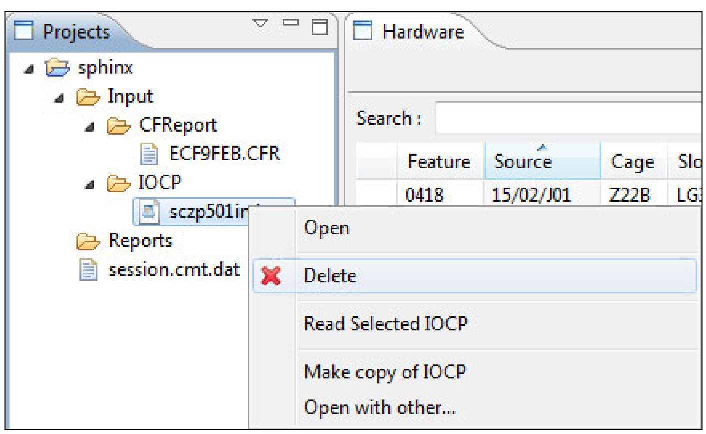

Figure 5-69 CMT - Reading the selected IOCP

A window displays the progress information (Figure 5-70).

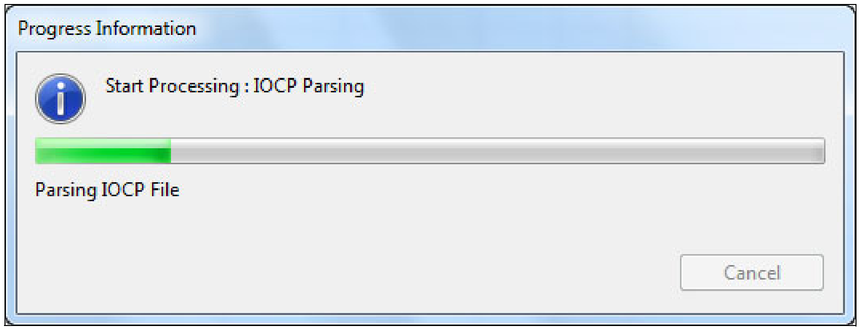

Figure 5-70 CMT - Processing the IOCP file

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 (Figure 5-71) includes three tabbed panes: Projects, Hardware Resolution, and Adapter Type Summary. Hardware Resolution is the middle pane and the Adapter Type Summary is on the right.

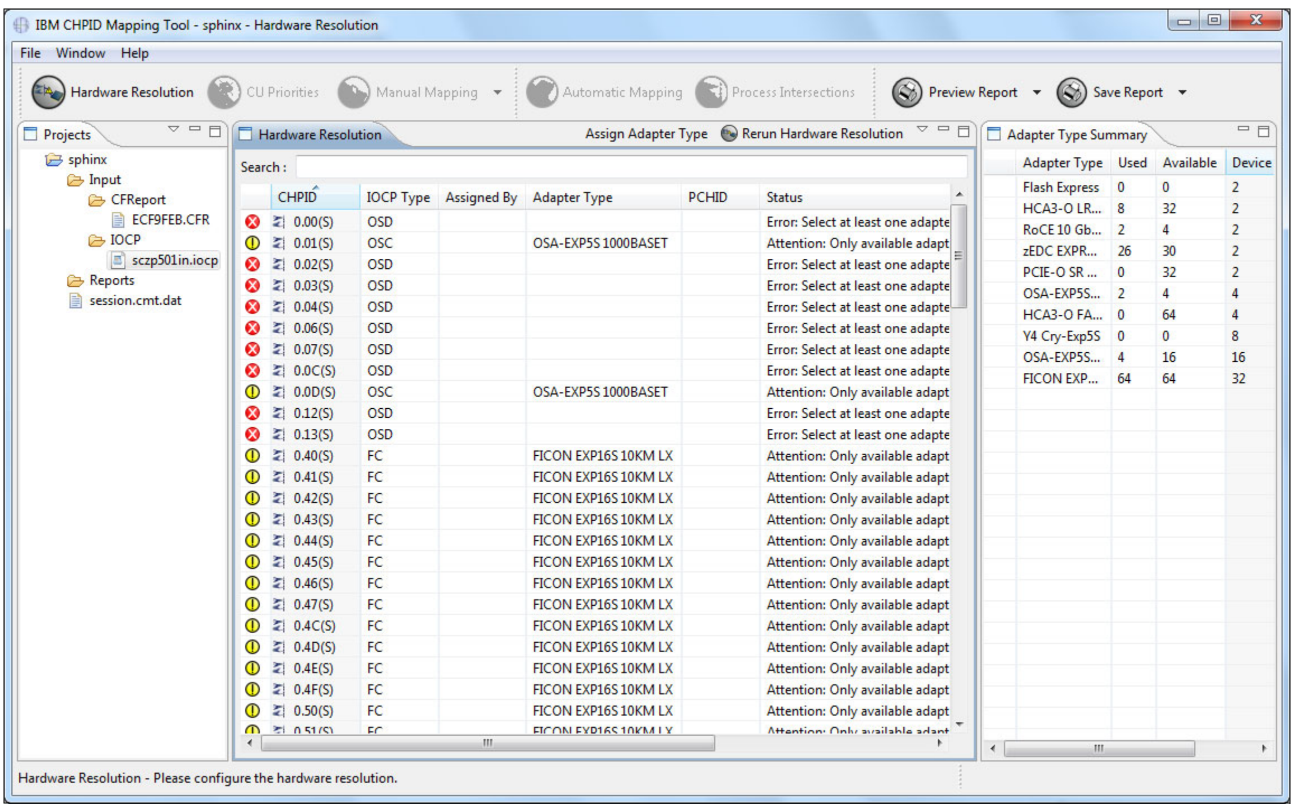

Figure 5-71 CMT - Hardware Resolution after Imported CFReport file

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 shows the following output:

- Hardware Resolution: This window lists all CHPIDs that were found; the Status column shows the CHPID information to be investigated. In the example, investigate the status; Table 5-3 lists status messages and possible resolutions.

Table 5-3 Status messages and possible resolutions

Status	Explanation	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 PCHID value for the old hardware. (The IOCP file contains a PCHID value for the hardware being removed.)	If you have any CHPIDs of IOCP type CIB, 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. Do the following steps for CIB CHPIDs: 1. Remove the AID values. 2. Do one of the following tasks: - Inside the CHPID Mapping Tool, perform manual mapping to associate these CHPIDs with AIDs. - Assign the AID values outside the tool, for example, using Hardware Configuration Definition (HCD). 3. Replace the IOCP file.
Select at least one adapter type.	A adapter type is not assigned to the current row.	Assign a adapter type to IOCP type.
Adapter_type is not compatible with *IOCP_type*.	Adapter type assigned for the CHPID is not compatible with the IOCP type specified by the IOCP file.	See Performing hardware resolution for a type mismatch
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 need to change the IOCP file or obtain additional hardware
PCHID_1 moved to new channel ID: PCHID_2 **Example**: 520 moved to 1E2	You are replacing hardware for an MES, and the IOCP file contains a PCHID value for the old hardware, which is being removed. This PCHID value has moved from an old machine to the PCHID value for the new hardware. PCHID_1 is the first PCHID value (for example, 520) and PCHID_2 is the second PCHID 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 if you prefer a different assignment.

- Manual mapping CIB CHPIDs: Availability Mapping cannot be used until all CIB CHPIDs are resolved. You can use manual mapping to resolve any CIB 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 PCHID in the Manual window. If this option is not selected (it has no check mark), then availability PCHIDs 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 might require recabling after availability assignments.

  Generally, select this option.

– `Reset CHPIDs assigned by CMT for config files`: The CFReport indicates that you are doing an MES/upgrade, and you have channels or CHPIDs (or both) that might have configuration files that are currently associated with them. The MES/upgrade might move some of those channel cards.

  Regardless of whether the channels are moving or not, the CHPID Mapping Tool either assigns PCHIDs 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 PCHID 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 only on CHPIDs that do not have PCHIDs assigned.

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

Attention: If you run **Reset CHPIDs assigned by IOCP**, it will reset any previously mapped CHPID assignments and 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 prior CHPID-to-PCHID relationship does not cause unacceptable availability.

5.6.3 Resolving CHPIDs with PCHID conflict

The CMT displays the CHPIDs with PCHID conflicts (Figure 5-72).

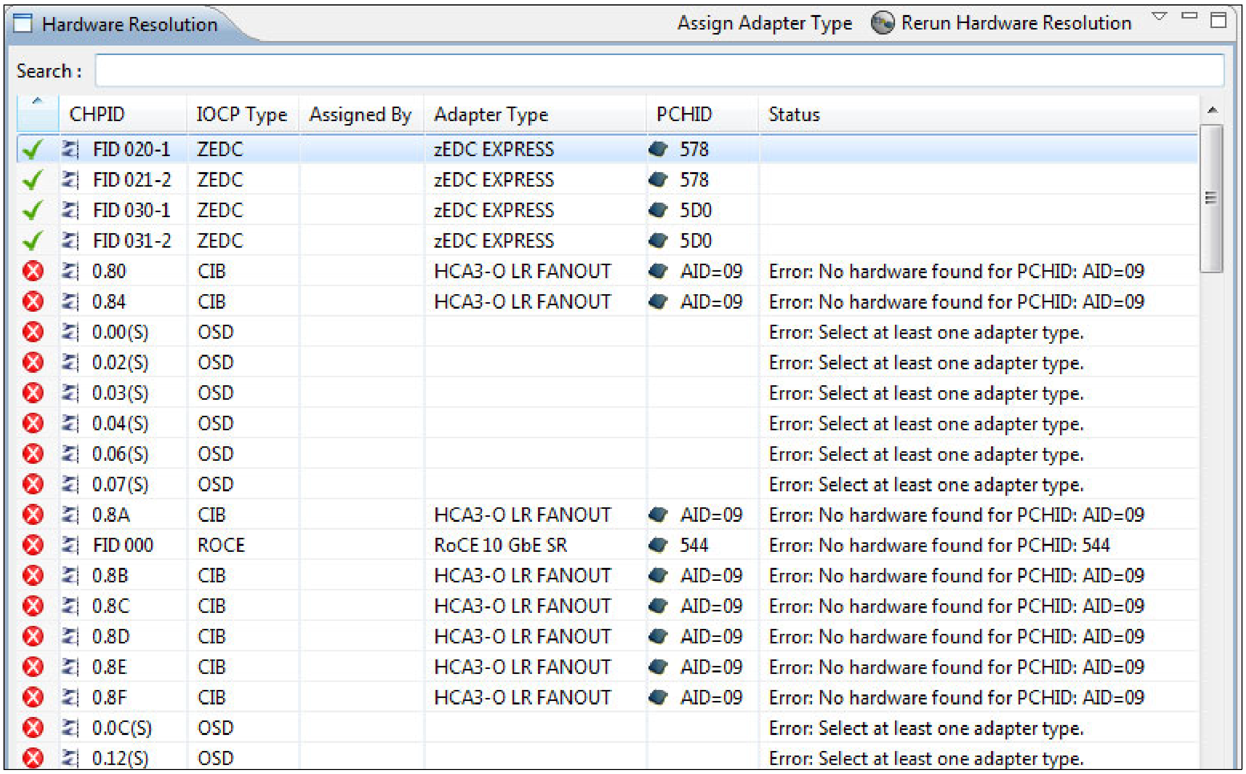

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 020-1	ZEDC		zEDC EXPRESS	578	
FID 021-2	ZEDC		zEDC EXPRESS	578	
FID 030-1	ZEDC		zEDC EXPRESS	5D0	
FID 031-2	ZEDC		zEDC EXPRESS	5D0	
0.80	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.84	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.00(S)	OSD				Error: Select at least one adapter type.
0.02(S)	OSD				Error: Select at least one adapter type.
0.03(S)	OSD				Error: Select at least one adapter type.
0.04(S)	OSD				Error: Select at least one adapter type.
0.06(S)	OSD				Error: Select at least one adapter type.
0.07(S)	OSD				Error: Select at least one adapter type.
0.8A	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
FID 000	ROCE		RoCE 10 GbE SR	544	Error: No hardware found for PCHID: 544
0.8B	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8C	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8D	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8E	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8F	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.0C(S)	OSD				Error: Select at least one adapter type.
0.12(S)	OSD				Error: Select at least one adapter type.

Figure 5-72 CMT - CHPIDs with PCHID conflicts

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

- An X in a red circle: This indicates an error.
- A green check mark: This indicates that the tool successfully resolved the specified Channel Type.

The example has these reasons to resolve hardware resolution issues:

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

5.6.4 Hardware resolution

In the example, the CHPID Mapping Tool displays an X in the first column of the Hardware Resolution pane (Figure 5-73) that is related to these error types: `No hardware found` and `FICON EXP8S 10KM LX is not compatible with OSD`.

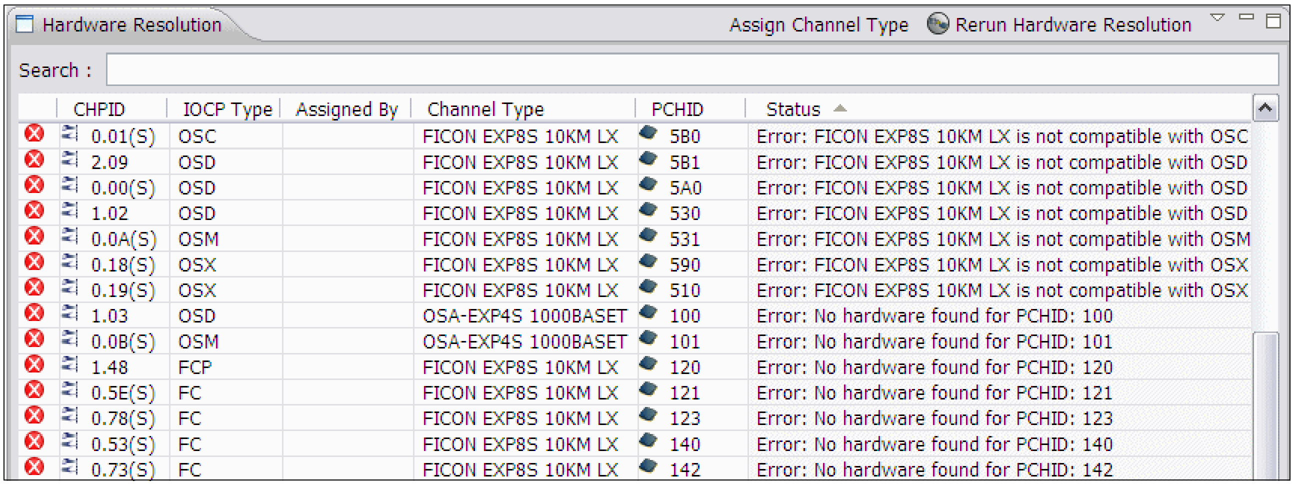
Hardware Resolution   Assign Channel Type   Rerun Hardware Resolution

Search :

CHPID	IOCP Type	Assigned By	Channel Type	PCHID	Status
0.01(S)	OSC		FICON EXP8S 10KM LX	5B0	Error: FICON EXP8S 10KM LX is not compatible with OSC
2.09	OSD		FICON EXP8S 10KM LX	5B1	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.00(S)	OSD		FICON EXP8S 10KM LX	5A0	Error: FICON EXP8S 10KM LX is not compatible with OSD
1.02	OSD		FICON EXP8S 10KM LX	530	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.0A(S)	OSM		FICON EXP8S 10KM LX	531	Error: FICON EXP8S 10KM LX is not compatible with OSM
0.18(S)	OSX		FICON EXP8S 10KM LX	590	Error: FICON EXP8S 10KM LX is not compatible with OSX
0.19(S)	OSX		FICON EXP8S 10KM LX	510	Error: FICON EXP8S 10KM LX is not compatible with OSX
1.03	OSD		OSA-EXP4S 1000BASET	100	Error: No hardware found for PCHID: 100
0.0B(S)	OSM		OSA-EXP4S 1000BASET	101	Error: No hardware found for PCHID: 101
1.48	FCP		FICON EXP8S 10KM LX	120	Error: No hardware found for PCHID: 120
0.5E(S)	FC		FICON EXP8S 10KM LX	121	Error: No hardware found for PCHID: 121
0.78(S)	FC		FICON EXP8S 10KM LX	123	Error: No hardware found for PCHID: 123
0.53(S)	FC		FICON EXP8S 10KM LX	140	Error: No hardware found for PCHID: 140
0.73(S)	FC		FICON EXP8S 10KM LX	142	Error: No hardware found for PCHID: 142

Figure 5-73 CMT - Hardware resolution status errors

More information: For more information about these error messages, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

The options that must be reset are as follows:

- Reset Incompatible (Hardware - I/O) Entries: (shown in example)
- Reset "Error: No hardware found" Entries: (shown in example)
- Reset "Select at least one adapter type": (shown in example)
- Reset "Required hardware for type IOCP_type not available": (*not* shown in example)
- Reset "PCHID_1 moved to new channel ID: PCHID_2": (*not* shown in example)

Reset Incompatible (Hardware - I/O) Entries

The Channel type that is assigned for the CHPID is not compatible with the IOCP type 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 PCHID. In the example, the IOCP type is OSD but the PCHID 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 (Figure 5-74).

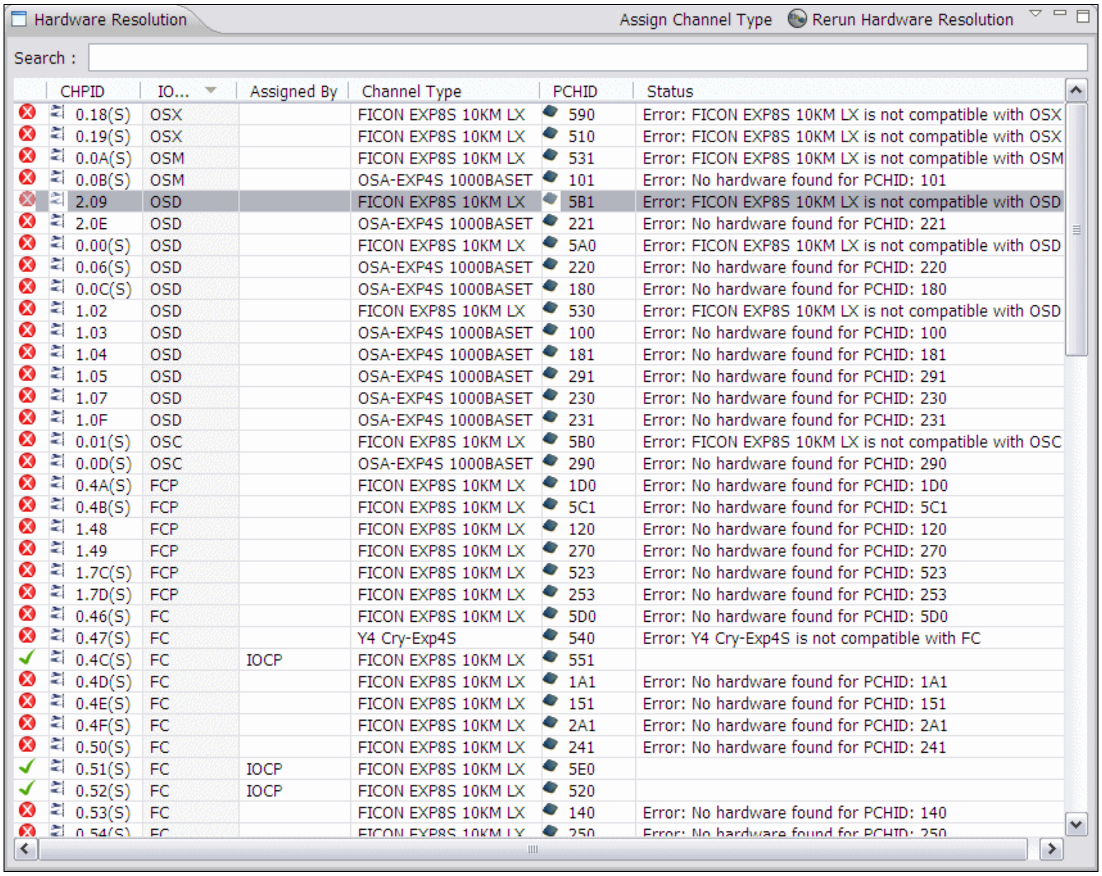

Figure 5-74   CMT - Channel_type is not compatible with IOCP_type

Select the channel type OSD. The Status is `Error: FICON EXP8S is not compatible with OSD`. Right-click in the row and select **Reset Incompatible (Hardware - I/O) Entries** to remove the PCHID values for only those rows (Figure 5-75).

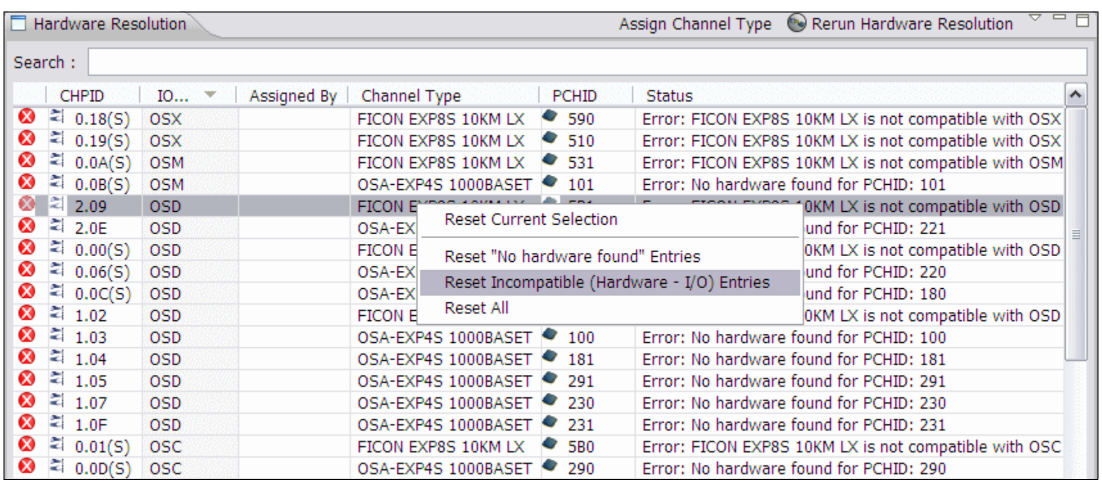

Figure 5-75   CMT - Channel_Type is not compatible with IOCP_type OSD

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 PCHIDs information (Figure 5-76).

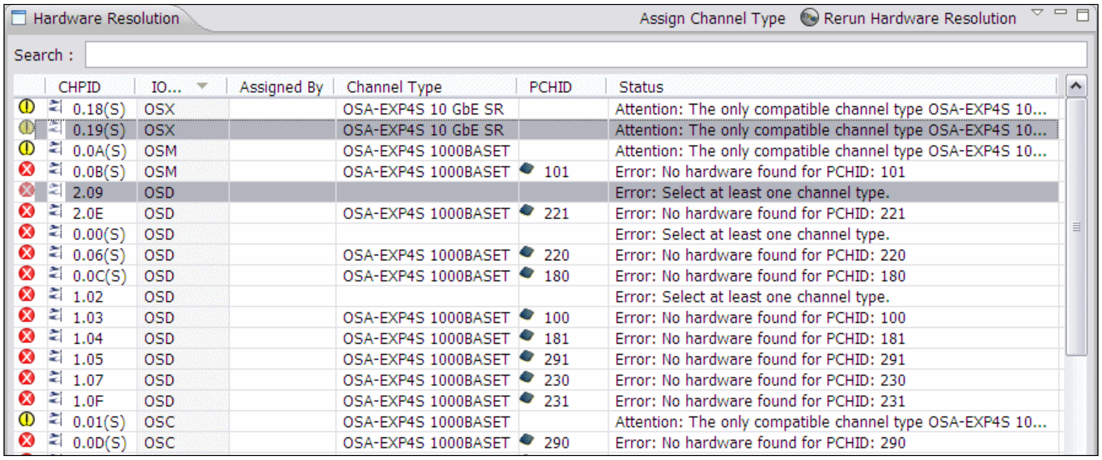

Figure 5-76 CMT - Results for reset of incompatible

The CHPID Mapping Tool now displays messages about any CHPID types that were imported from the IOCP input file (IODF) into the CMT that do not have any associated hardware support in the `CFReport` file (Figure 5-77). Click **OK**. The same figure also shows the Adapter Type Summary details.

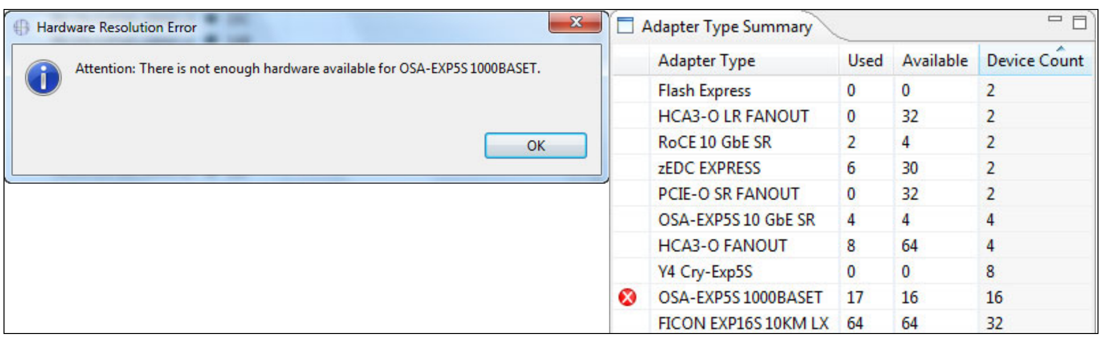

Figure 5-77 CMT - Required Hardware unavailable

Excessive numbers of *OSC* CHPID types are in the example IODF to show how the CHPID Mapping Tool handles this condition. For more information, see 5.2.8, “Overdefining channel paths on an XMP processor” on page 200.

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

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

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

To continue with this example, complete the following steps:

1. Return to the IODF and change the PCHID values for the OSD CHPIDs (or any other CHPIDs that have no 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 panel and selecting **Import IOCP File**.

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

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

Reset "Error: No hardware found" Entries

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` (Figure 5-78).

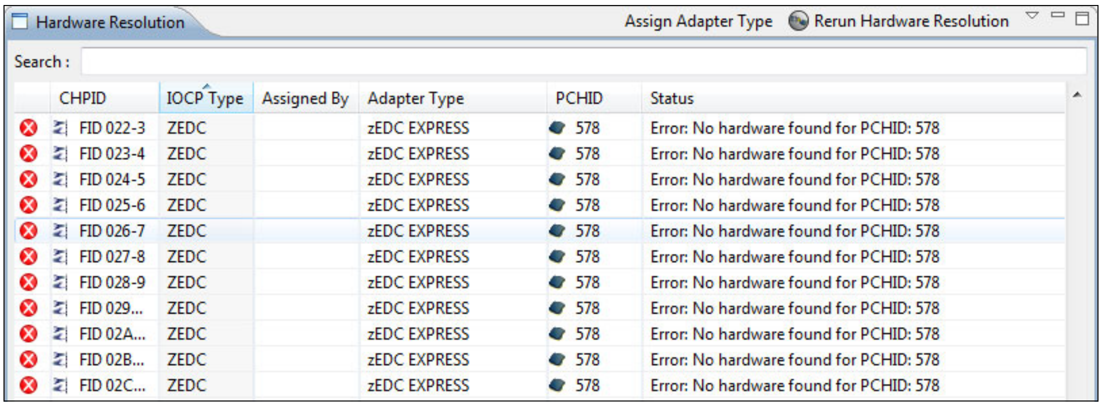

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 022-3	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 023-4	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 024-5	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 025-6	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 026-7	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 027-8	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 028-9	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 029...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02A...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02B...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02C...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578

Figure 5-78 CMT - 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 PCHID values for those rows (Figure 5-79).

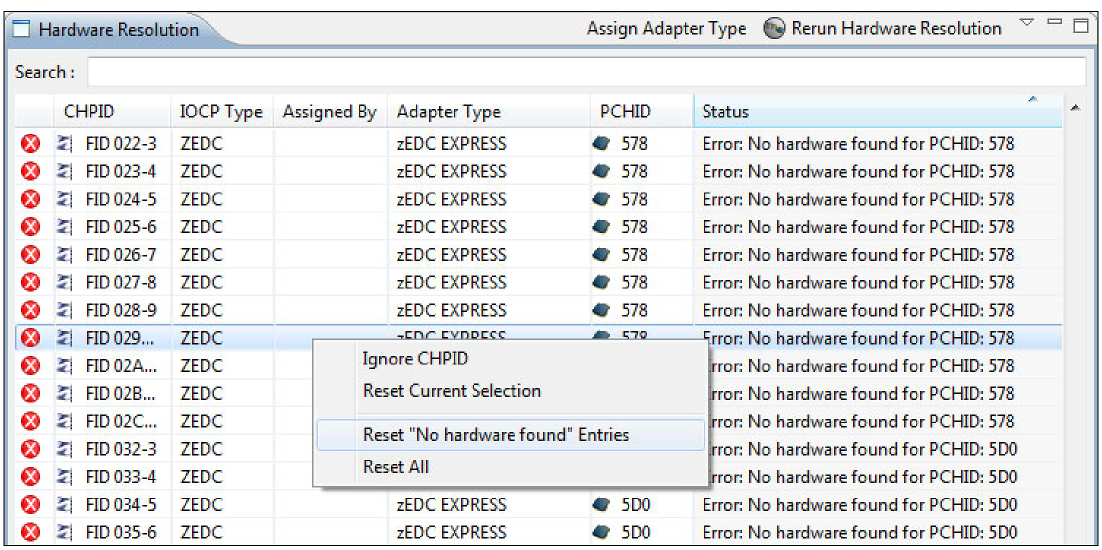

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 022-3	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 023-4	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 024-5	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 025-6	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 026-7	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 027-8	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 028-9	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 029...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02A...	ZEDC				rror: No hardware found for PCHID: 578
FID 02B...	ZEDC				rror: No hardware found for PCHID: 578
FID 02C...	ZEDC				rror: No hardware found for PCHID: 578
FID 032-3	ZEDC				rror: No hardware found for PCHID: 5D0
FID 033-4	ZEDC				rror: No hardware found for PCHID: 5D0
FID 034-5	ZEDC		zEDC EXPRESS	5D0	Error: No hardware found for PCHID: 5D0
FID 035-6	ZEDC		zEDC EXPRESS	5D0	Error: No hardware found for PCHID: 5D0

Figure 5-79 CMT - Resetting No Hardware found entries

The tool replaces the X with an *Attention* icon, changes the status message, and removes the PCHID information (Figure 5-80).

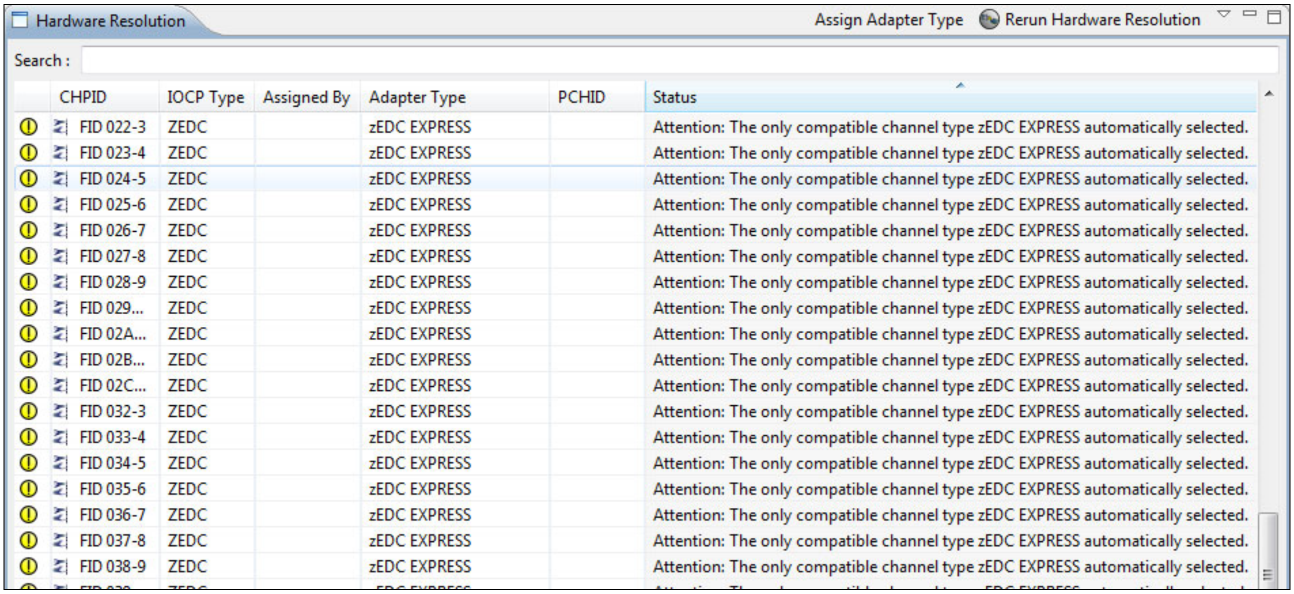

Figure 5-80 CMT - Results of resetting No hardware found

Reset “Select at least one adapter type”

The adapter type is not assigned to the current row. Assign an adapter type to the IOCP type:

1. Click the **Adapter Type** column in the target row. The tool displays an arrow in the Channel Type column of the target row (Figure 5-81).

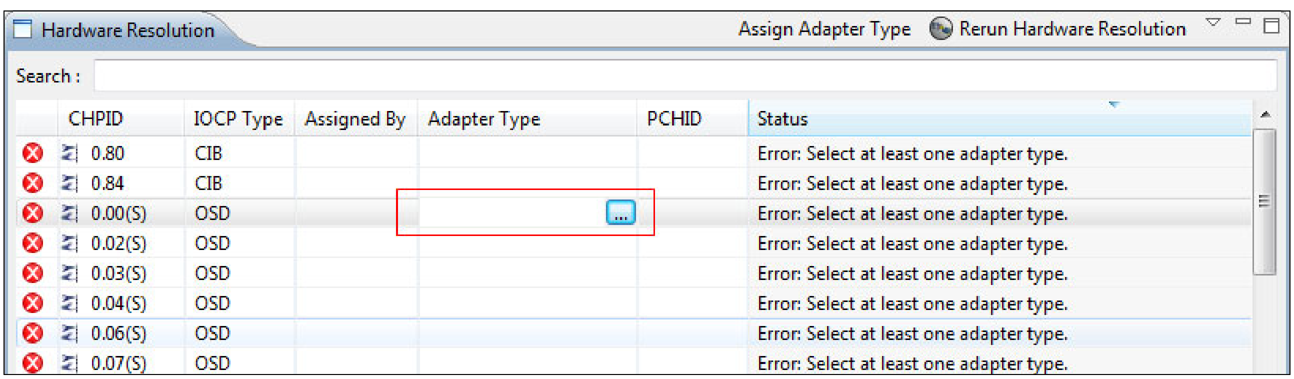

Figure 5-81 CMT - Selecting at least one adapter type

2. Click the ellipses (...) box.

3. The tool displays a list of available and compatible card types for the CHPID as shown in Figure 5-82. Select an adapter type and click **OK**.
4. In the Adapter Type Summary tab, observe that the “Used” and “Available” totals change.

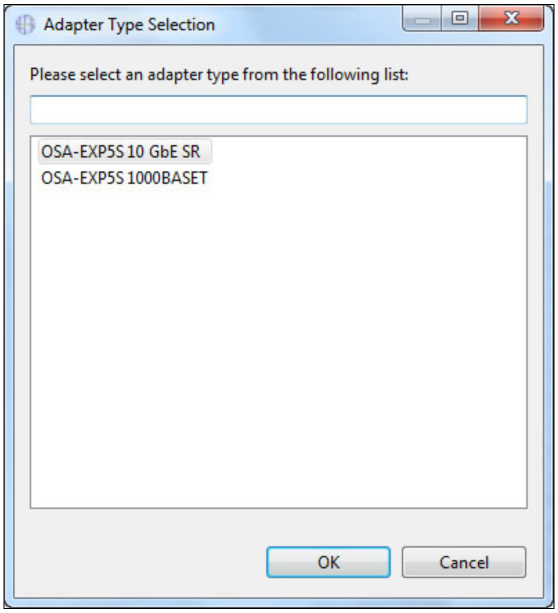

Figure 5-82 CMT - Adapter Type Selection

Reset “Required hardware for type IOCP_type not available”

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

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

You must change IOCP file or obtain more hardware.

Reset “PCHID_1 moved to new channel ID: PCHID_2”

When moving from old hardware to new hardware, for example during a miscellaneous equipment specification (MES), the PCHID value assigned to a feature can change. This message indicates that the IOCP file contains a PCHID value for the old machine that is being removed. The PCHID value is changed from the old machine to the PCHID value for the new machine. For example, PCHID_1 is the first PCHID value representing the old hardware (for example, 1B0) and PCHID_2 is the new value representing the new hardware (for example, 533). In essence, the feature is present in both the old and new hardware, but its location (PCHID) has 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** button becomes available.

5.6.5 Manual mapping to resolve CIB CHPIDs

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

To resolve the CIB CHPIDs, assign the available CHPIDs. Click **Manual Mapping** (Figure 5-83).

Figure 5-83 CMT - Manual Mapping

Ensure that the tool is set to display Manual Mapping in the **Hardware -> I/O** format (Figure 5-84).

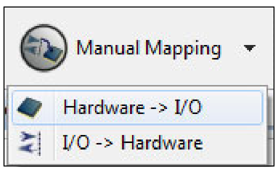

Figure 5-84 CMT - Manual Mapping of Hardware -> I/O

Click every row that has type HCA3-O in the Channel Type column. The tool displays all the available CHPIDs with IOCP type (Figure 5-85).

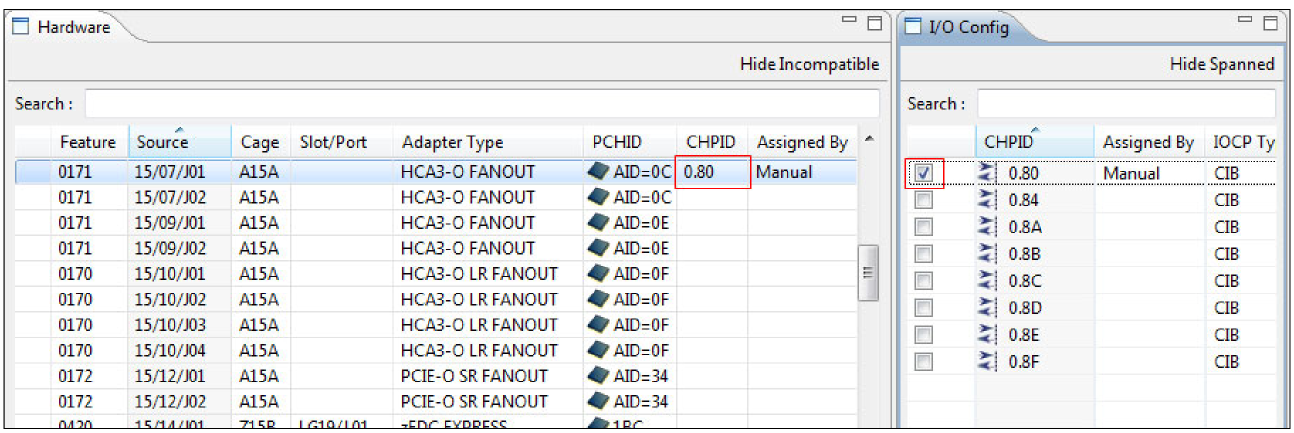

Figure 5-85 CMT - Adapter Type of HCA3 and associated CHPID assigned

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 HCA3-O adapter type, you see the `Multiple -->` value (Figure 5-86 on page 223) inserted in the CHPID and Assigned By columns.

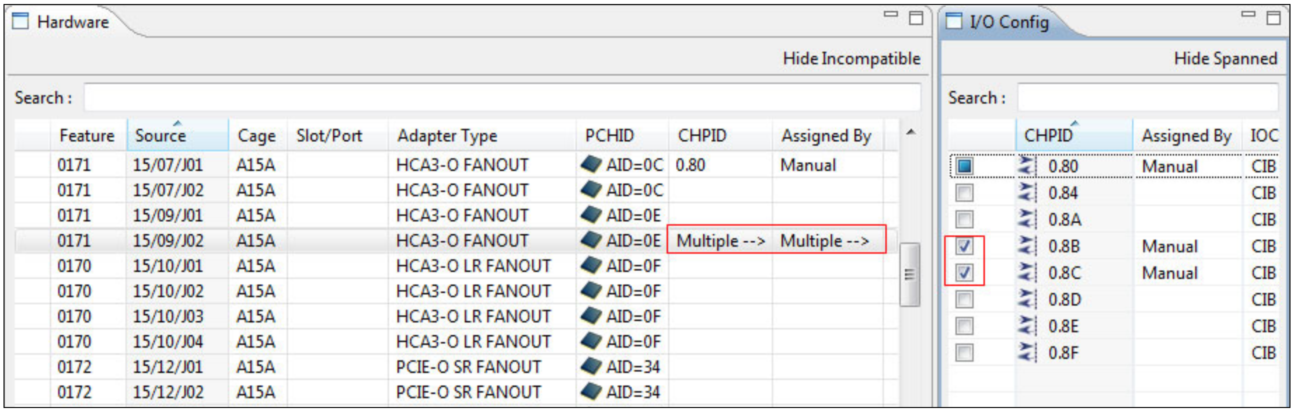

Figure 5-86 CMT - Adapter Type of HCA3 and associated multiple CHPID assigned

The **Automatic Mapping** button becomes available after you assign all the CHPIDs of IOCP type CIB.

5.6.6 Processing the CU Priority and Automatic Mapping

If you are importing an IOCP statements file from a 2817 or 2827 that had CU Priority values defined, review the CU Priority values first. The CHPID Mapping Tool can then perform the availability functions appropriately for a 2964.

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 these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices (Figure 5-87 on page 224). For the example, select the following two options and then click **OK**:
   - **Reset CHPIDs assigned by Automatic Mapping**
   - **Reset CHPIDs assigned by IOCP**

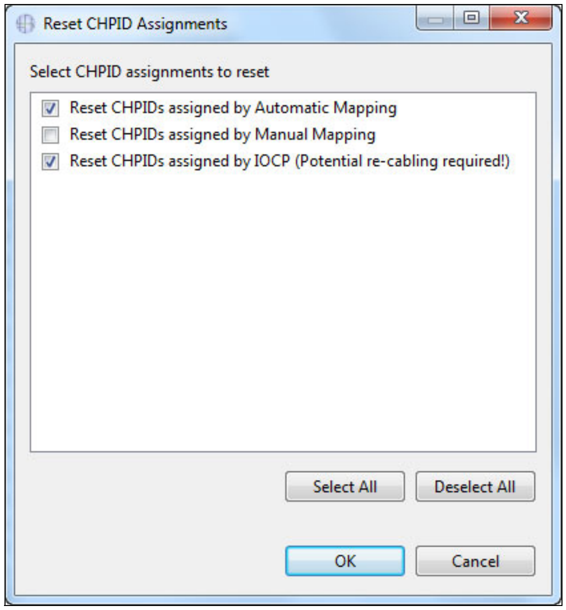

Figure 5-87 CMT - Reset CHPID Assignments

> **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 (Figure 5-88).

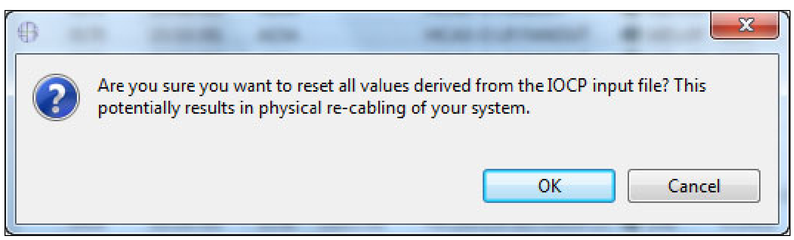

Figure 5-88 CMT - Reset CHPID assignments warning message

4. The 2964 has availability rules that differ from 2817 and 2827, so remove all PCHID 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 (Figure 5-89). Click **OK**.

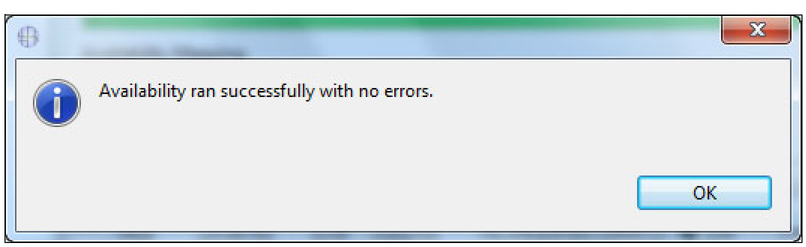

Figure 5-89 CMT - Availability ran successfully with no errors message

7. Click **OK** (Figure 5-90).

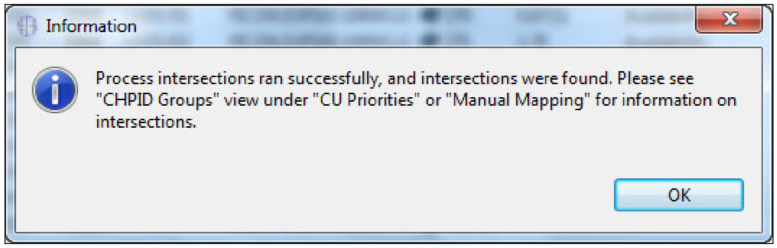

Figure 5-90 CMT - Process Intersections run successfully message

The possible intersects are as follows:

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 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 if they are acceptable (Figure 5-91). The example returned the “B” intersect.

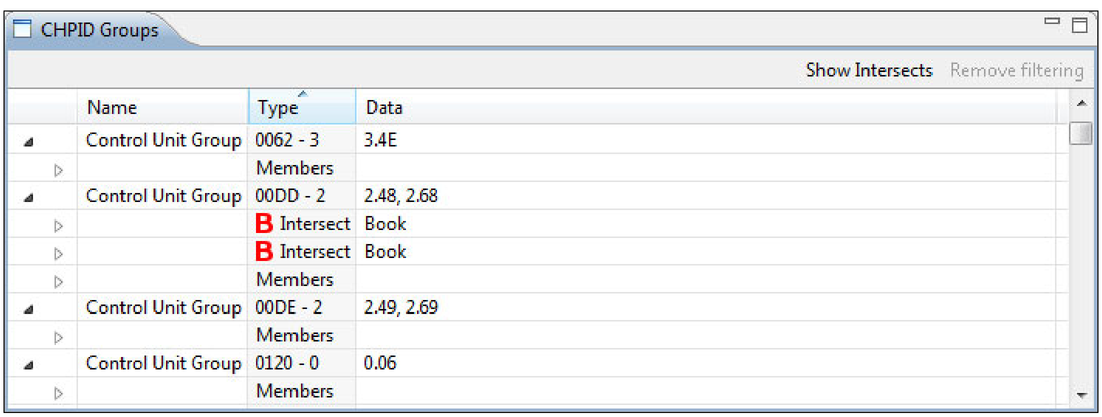

Figure 5-91 CMT - B Intersect examples

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.

Check and set values for items such as OSA-ICC CHPIDs and FCTC CHPIDs to ensure that the CHPID Mapping Tool allocates these CHPIDs with high PCHID availability.

1. Click **CU Priorities**. By default, this pane is in the center at the top.
2. In the CU Priorities pane, search in the CU Number column for the control units that you want to set a priority for.

3. Type 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. In the example, set the OSC type CU Numbers to priority 333 (Figure 5-92).

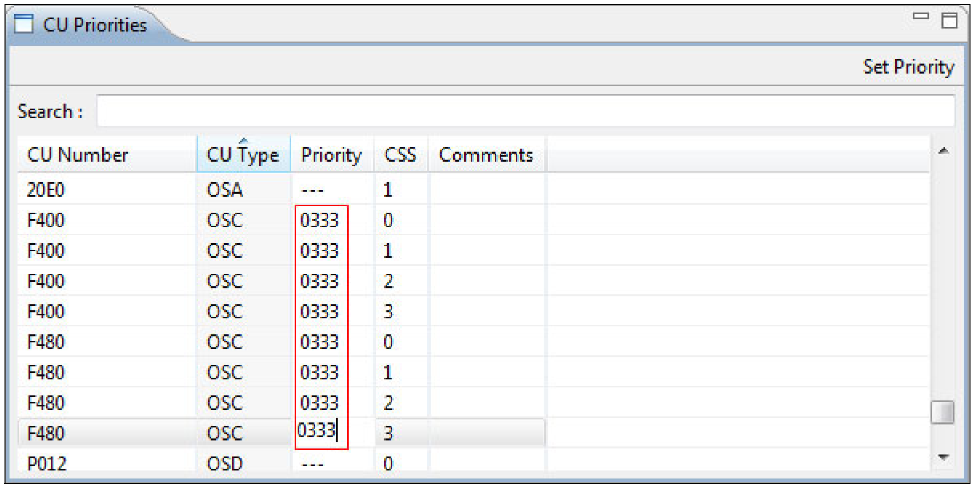

Figure 5-92 CMT - Set CU Priorities

If coupling links are used by a CF image, group those links.

Group each set of CHPIDs going to a different CPC with a common priority. For example, suppose the CF image has four links (CHPIDs 40, 41, 42, and 43) and that 40 and 41 go to one CPC, and 42 and 43 go to a different CPC. In this case, give CHPIDs 40 and 41 one priority and CHPIDs 42 and 43 a different priority. The concept is the same regardless of the number of connecting CPCs or the number of links to each CPC.

Now perform the second availability function by completing these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices. Click **Reset CHPIDs assigned by Automatic Mapping**.
3. Click **OK**.

   The Hardware Resolution pane shows that the CHPID and Assigned By columns are no longer blank (Figure 5-93). The CMT assigned CHPIDs to PCHIDs and placed the Availability value in the Assigned By column, indicating that the CHPID values were assigned based on availability.

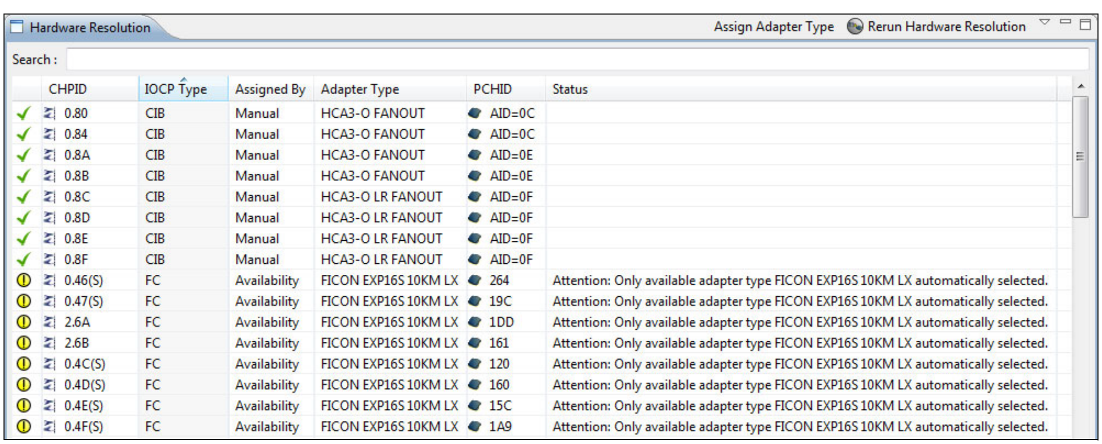

Figure 5-93 CMT - CHPIDs assigned

The possible Assigned By column values are as follows:

Manual You made the assignment by using manual mapping.

Automatic You made the assignment by using automatic mapping.

IOCP The IOCP source made the assignment.

Config File The CHPID Mapping Tool forced an assignment because of configuration file requirements.

More information: See the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

4. You can now display the results of the channel mapping. You can also sort the report in various ways. For example, to see how the CHPID Mapping Tool ranked the control units, select the **CU Priorities** pane and click the **Priority** column (Figure 5-94).

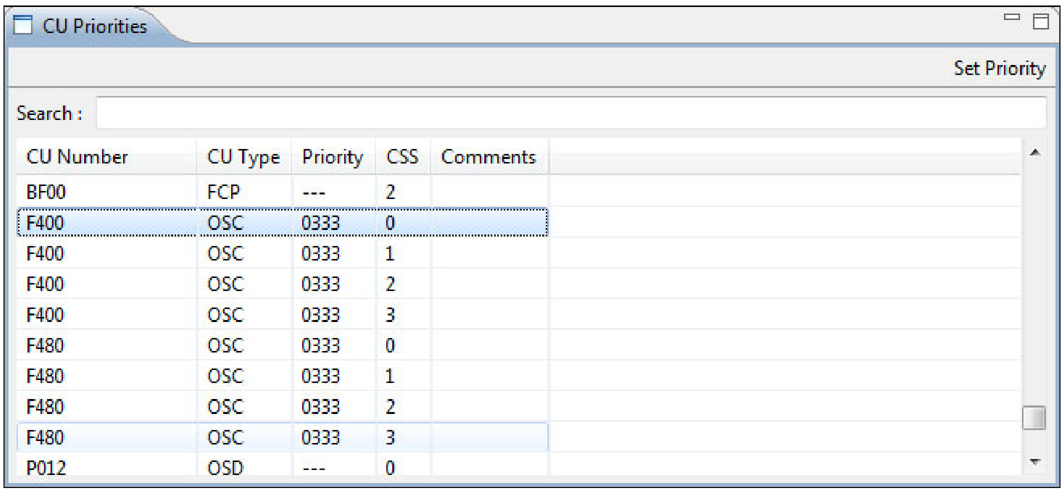

Figure 5-94 CMT - CU Priorities showing assigned priorities

The example illustrates how CU Priority values are represented in the IOCP file.

Tip: The control unit priorities are stored in the IOCP output file created by the CMT that gets migrated back into HCD. HCD maintains these priorities and outputs them when it creates another IOCP deck. They are in the form of commented lines at the end of the IOCP deck, as shown here:

```
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
```

5.6.7 CHPIDs not connected to control units

In the CU Priorities window, click in the **CU Number** column (Figure 5-95). The CHPID Mapping Tool shows, at the end of the list, all CHPIDs 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.

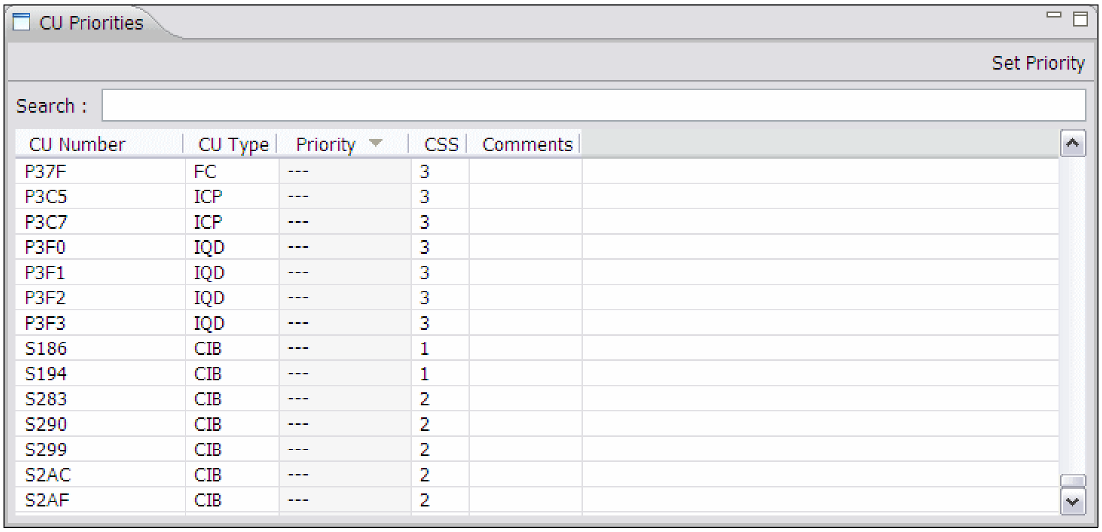

CU Number	CU Type	Priority	CSS	Comments
P37F	FC	---	3	
P3C5	ICP	---	3	
P3C7	ICP	---	3	
P3F0	IQD	---	3	
P3F1	IQD	---	3	
P3F2	IQD	---	3	
P3F3	IQD	---	3	
S186	CIB	---	1	
S194	CIB	---	1	
S283	CIB	---	2	
S290	CIB	---	2	
S299	CIB	---	2	
S2AC	CIB	---	2	
S2AF	CIB	---	2	

Figure 5-95 CMT - 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 need to 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 there are extra CHPIDs for anticipated new control units, consider grouping these CHPIDs with a common priority. Having a common priority allows the availability mapping function to pick PCHIDs that can afford your new control unit availability.

5.6.8 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**. Figure 5-96 shows the options to create a Preview Report or Save Report.

Figure 5-96 CMT - 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 Save Report lists an extra selection (Figure 5-97).

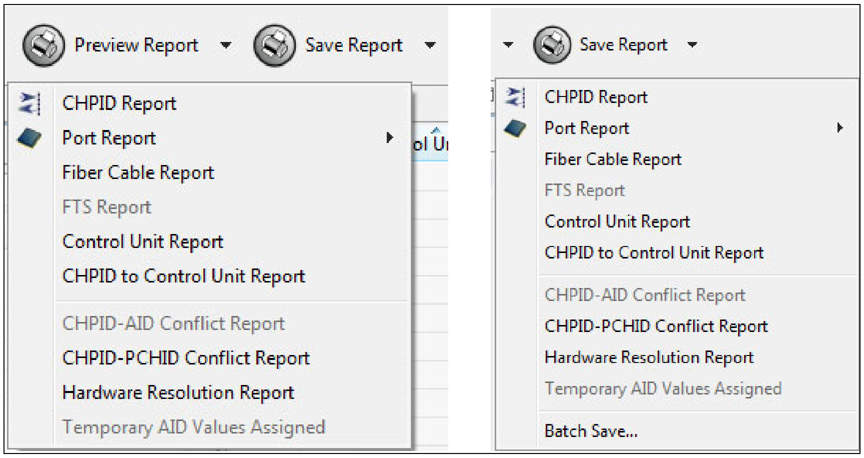

Figure 5-97 CMT - Preview Report and Save Report menus

For simplicity, only three reports are described in this example: The CHPID Report; the Port Report, sorted by location; and the 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 PCHID or cage/slot/port information before system delivery.

CHPID Report

To create the CHPID report, complete the following steps:

1. Click **Preview Report** → **CHPID Report** (Figure 5-98).

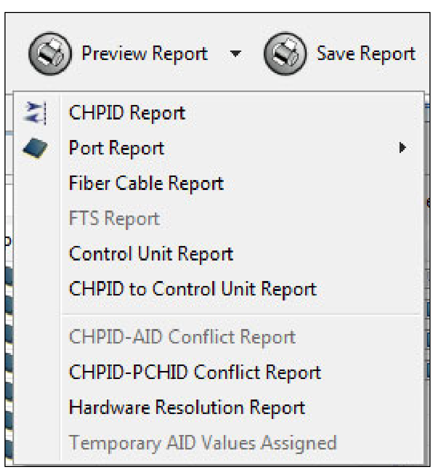

Figure 5-98 CMT - Preview report - CHPID Report

The CHPID Mapping Tool displays the CHPID Report in a Report tab within the CMT (Figure 5-99).

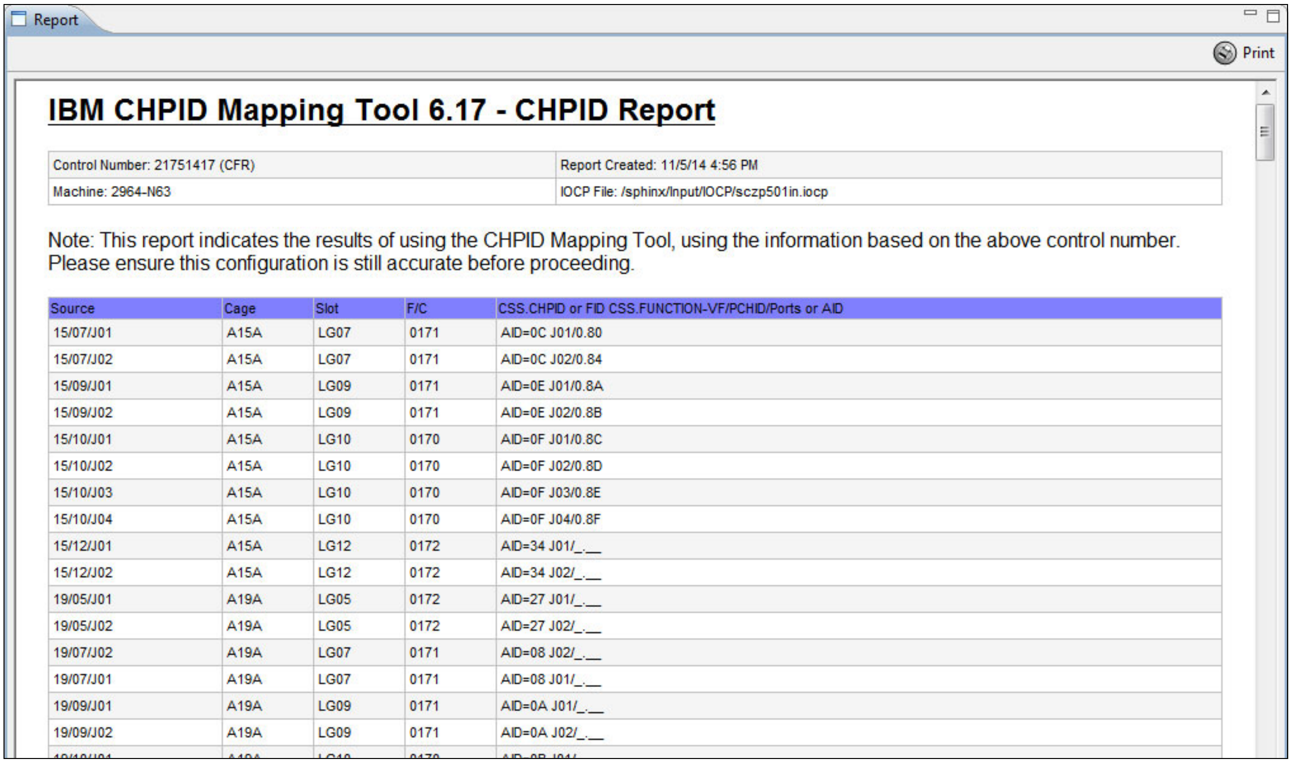

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 4:56 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__

Figure 5-99 CMT - 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 (Figure 5-100). 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**.

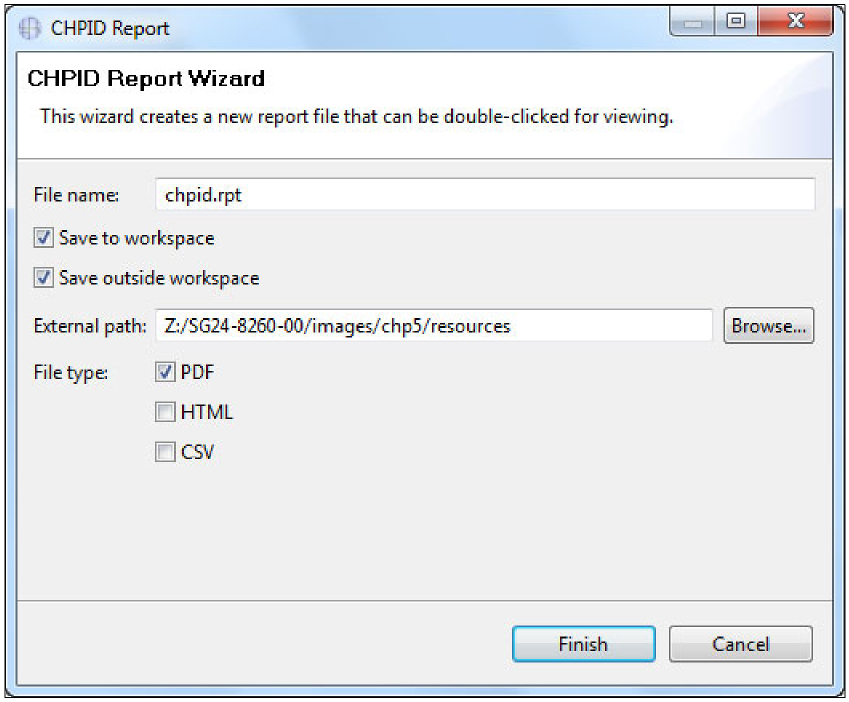

Figure 5-100 CMT - Save CHPID Report

The CHPID Report is created by the CHPID Mapping Tool (Figure 5-101).

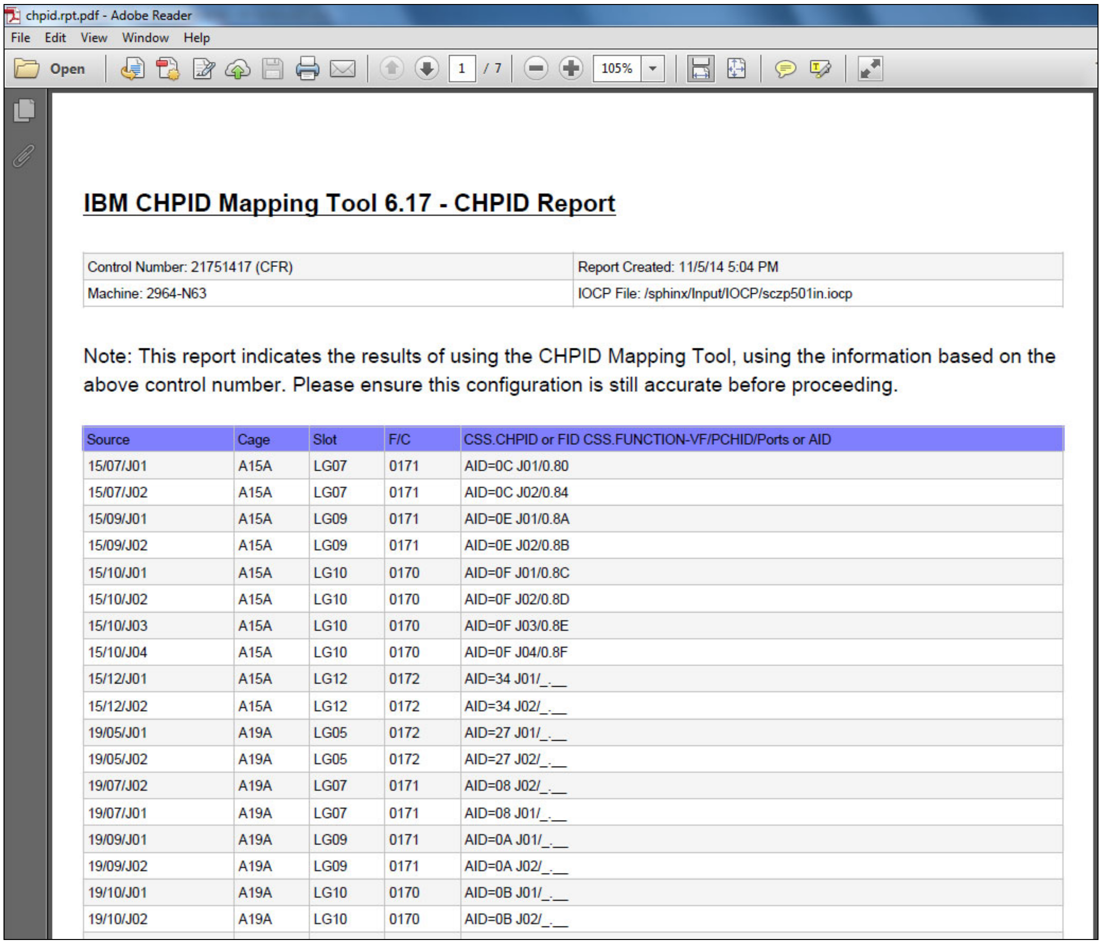

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:04 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__
19/10/J01	A19A	LG10	0170	AID=0B J01/_.__
19/10/J02	A19A	LG10	0170	AID=0B J02/_.__

Figure 5-101 CMT - CHPID Report example in PDF format

At the end of this CHPID Report is a list of CHPIDs with modified PCHID/AID assignments (Figure 5-102). This report is valuable for moving cables.

List of CHPIDs having modified PCHID/AID assignments

Note: For CHPIDs that had PCHID/AID assignments in the IOCP file that was loaded for this session of the Mapping Tool.

CHPIDs or FUNCTIONs	Previous PCHID/AID-Port	PCHID/AID-Port	Current Location	F/C
0.80	09-1	0C-1	A15ALG07J.01	0171
0.84	09-1	0C-02	A15ALG07J.02	0171
0.8A	09-1	0E-1	A15ALG09J.01	0171
0.8B	09-1	0E-02	A15ALG09J.02	0171
0.8C	09-1	0F-01	A15ALG10J.01	0170
0.8D	09-1	0F-02	A15ALG10J.02	0170
0.8E	09-1	0F-03	A15ALG10J.03	0170
0.8F	09-1	0F-04	A15ALG10J.04	0170
FID 000	544	208	Z08BLG03J.01	0411
FID 010	5EC	140	Z22BLG20J.01	0411
FID 020-1	578	1BC	Z15BLG19J.01	0420
FID 021-2	578	27C	Z08BLG38J.01	0420
FID 022-3	578	Not Assigned		
FID 023-4	578	Not Assigned		

Figure 5-102 CMT - List of CHPIDs that have modified PCHID/AID assignments

CHPID to Port Report, sorted by location

To create the Port Report, 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 (Figure 5-103).

IBM CHPID Mapping Tool 6.17 - CHPID to Port Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:12 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
A15A	07/LG	AID=0C / J.01	15/07	HCA3-O FANOUT	0.80	Manual
A15A	07/LG	AID=0C / J.02	15/07	HCA3-O FANOUT	0.84	Manual
A15A	09/LG	AID=0E / J.01	15/09	HCA3-O FANOUT	0.8A	Manual
A15A	09/LG	AID=0E / J.02	15/09	HCA3-O FANOUT	0.8B	Manual
A15A	10/LG	AID=0F / J.01	15/10	HCA3-O LR FANOUT	0.8C	Manual
A15A	10/LG	AID=0F / J.02	15/10	HCA3-O LR FANOUT	0.8D	Manual
A15A	10/LG	AID=0F / J.03	15/10	HCA3-O LR FANOUT	0.8E	Manual
A15A	10/LG	AID=0F / J.04	15/10	HCA3-O LR FANOUT	0.8F	Manual
A15A	12/LG	AID=34 / J.01	15/12	PCIE-O SR FANOUT		
A15A	12/LG	AID=34 / J.02	15/12	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.01	19/05	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.02	19/05	PCIE-O SR FANOUT		
A19A	07/LG	AID=08 / J.01	19/07	HCA3-O FANOUT		
A19A	07/LG	AID=08 / J.02	19/07	HCA3-O FANOUT		
A19A	09/LG	AID=0A / J.01	19/09	HCA3-O FANOUT		

Figure 5-103 CMT - CHPID to Port Report, sorted by location

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 (Figure 5-104).

IBM CHPID Mapping Tool 6.17 - CHPID to CU Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:15 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
0	00	OSD	19/12/J01	Z08B LG07 J00J01	214	2040	OSA	---
0	01	OSC	19/03/J01	Z15B LG21 J00J01	1C4	F400	OSC	0333
0	02	OSD	15/03/J01	Z15B LG35 J00J01	1F0	2080	OSA	---
0	03	OSD	15/15/J01	Z22B LG17 J00J01	134	20A0	OSA	---
0	04	OSD	15/02/J01	Z22B LG36 J00J01	174	20C0	OSA	---
0	06	OSD	19/02/J01	Z22B LG21 J00J01	144	0120	OSA	---
0	07	OSD	15/02/J01	Z22B LG37 J00J01	178	0140	OSA	---
0	0C	OSD	19/03/J01	Z15B LG22 J00J01	1C8	2060	OSA	---
0	0D	OSC	15/15/J01	Z22B LG18 J00J01	138	F480	OSC	0333
0	12	OSD	15/05/J01	Z08B LG36 J00J01	274	P012	OSD	---
0	13	OSD	19/14/J01	Z15B LG06 J00J01	190	P013	OSD	---
0	40	FC	15/14/J01	Z15B LG12 D1	1A4	1000	2107	---
						1200	2107	---
						1400	2107	---

Figure 5-104 CMT - CHPID to CU Report

5.6.9 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 has PCHIDs that are assigned to CHPIDs.

Using HCM: You might prefer to use HCM to transfer the updated IOCP statements file back to the host. However, first run the next step in the CHPID Mapping Tool to create the updated IOCP file.

To create the IOCP, complete the following steps:

1. Select **File** → **Export IOCP input file** (Figure 5-105).

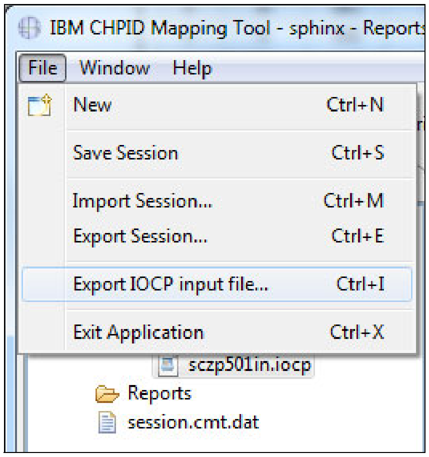

Figure 5-105 CMT - Export IOCP input file

2. Enter the Export Path and IOCP Name for the IOCP output file and click **Finish** (Figure 5-106).

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

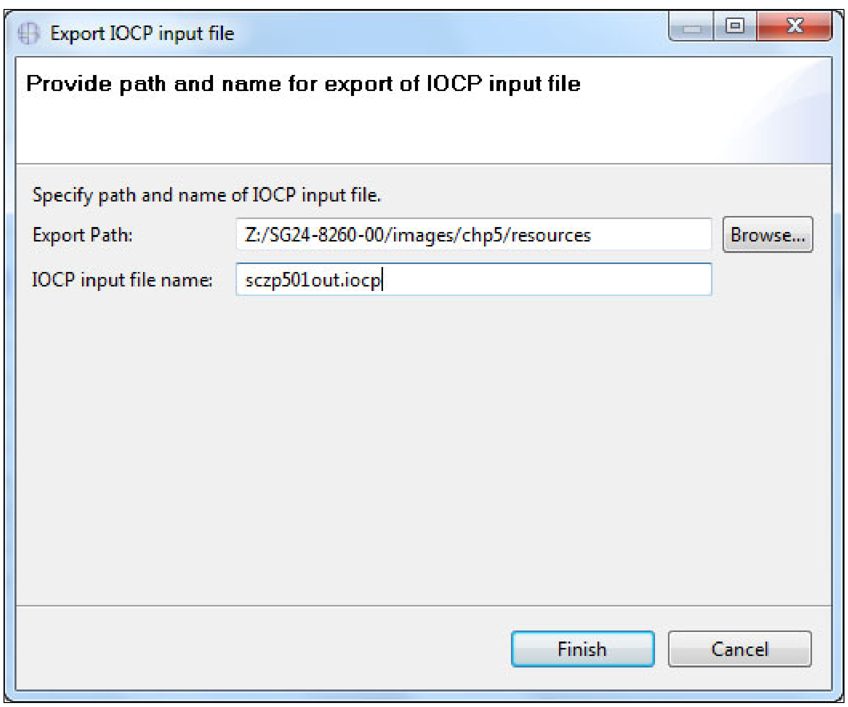

Figure 5-106 CMT - Export IOCP File

3. Select **File** → **Save Session** (Figure 5-107).

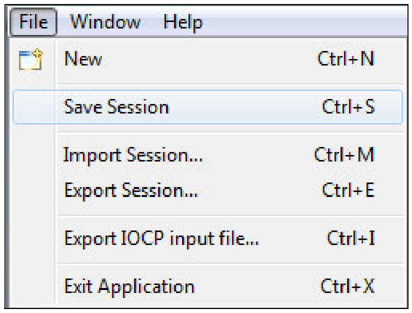

Figure 5-107 CMT - Save session

You might want to save your project before exiting the CMT application.

5.7 HCD: Updating the 2964 work IODF with PCHIDs

After you map the PCHIDs to CHPIDs by using the CHPID Mapping Tool, transfer this information back into HCD. To update the IODF with the PCHIDs, follow these following steps:

1. Upload the IOCP file that was created by the CMT (`SCZP501out.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 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. They are in the form of commented lines at the end of the IOCP deck (Example 5-6).

Example 5-6 HCD - Updated IOCP statements file: with CMT statements

```
         CNTLUNIT CUNUMBR=FFFD,PATH=((CSS(0),C6,C7)),UNIT=CFP
         IODEVICE ADDRESS=(FFCF,007),CUNUMBR=(FFFD),UNIT=CFP
         IODEVICE ADDRESS=(FFDD,007),CUNUMBR=(FFFD),UNIT=CFP
         CNTLUNIT CUNUMBR=FFFE,PATH=((CSS(0),C4,C5)),UNIT=CFP
         IODEVICE ADDRESS=(FFD6,007),CUNUMBR=(FFFE),UNIT=CFP
         IODEVICE ADDRESS=(FFF2,007),CUNUMBR=(FFFE),UNIT=CFP
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
******************************** Bottom of Data ************************
```

2. From the HCD main panel, enter the work IODF name used. Select option **5. Migrate configuration data** (Figure 5-108).

```
                           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 . . . 'SYS6.IODFB5.WORK'                     +
```

Figure 5-108 HCD - Hardware Configuration: migrate configuration data

3. From the Migrate Configuration Data panel (Figure 5-109), 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
```

Figure 5-109 HCD - Migrate Configuration Data: migrate IOCP/OS data.

4. The Migrate IOCP Data panel opens (Figure 5-110 on page 238). Compete 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 OS configuration that is associated with the processor.

IOCP only input data set This data set is the one specified when the `SCZP501out.iocp` file was uploaded to z/OS.

Processing mode Select option **2** to save the results of the migration. (Before using option 2, however, 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 . . . . . . . . . . . . SCZP501   +   CSS ID . . . . . . _  +
 OS configuration ID  . . . . . . . . L06RMVS1  +

 Combined IOCP/MVSCP input data set . _____________________________________
 IOCP only input data set . . . . . . 'SYS6.IODFB5.IOCPOUT.SCZP501'
 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)
```

Figure 5-110 HCD - 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 (Figure 5-111).

```
------------------------- 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
#                    SYS6.IODFB5.WORK.
***************************** Bottom of data ******************************
```

Figure 5-111 HCD - Migration Message List: successful message

The work IODF now contains both the CHPID definitions and the mapping to PCHIDs 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.

5.8 HCD: Building the 2964 production IODF

To use the definitions that were updated in HCD, create a 2964 production IODF from your work IODF. Then, remotely or locally write the IODF to the 2964 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** (Figure 5-112).

```
                           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.IODFB5.WORK'                     +
```

Figure 5-112 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel is displayed (Figure 5-113). 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
```

Figure 5-113 HCD - Activate or Process Configuration Data: build production IODF

3. HCD displays the Message List panel (Figure 5-114). Verify that you have only severity “W” (warning) messages and that they are normal for the configuration. Correct any other messages that should not occur and try to build the production IODF again. Continue this process until you have no messages that indicate problems.

```
------------------------------- Message List -------------------------------
   Save  Query  Help
 --------------------------------------------------------------------------
                                                                 Row 1 of 18
 Command ===> ___________________________________________ Scroll ===> CSR

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

 / Sev Msg. ID  Message Text
 _ W   CBDA857I No channel paths attached to partition A0F of processor
 #              SCZP501.0.
 _ W   CBDA857I No channel paths attached to partition A1B of processor
 #              SCZP501.1.
 _ W   CBDA857I No channel paths attached to partition A1D of processor
 #              SCZP501.1.
 _ W   CBDA857I No channel paths attached to partition A1E of processor
 #              SCZP501.1.
 _ W   CBDA857I No channel paths attached to partition A1F of processor
 #              SCZP501.1.
 _ W   CBDA857I No channel paths attached to partition A3B of processor
```

Figure 5-114 HCD - Message List: building production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (Figure 5-115). Complete the Production IODF name and Volume serial number fields, and then press Enter.

```
-------------- Build Production I/O Definition File -----------


 Specify the following values, and choose how to continue.

 Work IODF name . . . : 'SYS6.IODFB5.WORK'

 Production IODF name . 'SYS6.IODFB5'______________________
 Volume serial number . IODFPK  +

 Continue using as current IODF:
 2   1.  The work IODF in use at present
     2.  The new production IODF specified above
```

Figure 5-115 HCD - Build Production I/O Definition File: data fields to be updated

6. The Define Descriptor Fields panel opens (Figure 5-116). 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.IODFB5'

Descriptor field 1  . . . SYS6
Descriptor field 2  . . . IODFB5
```

Figure 5-116 HCD - 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.IODFB5 created.`

Proceed to the next section to implement the configuration on the 2827 in preparation for its upgrade to a 2964.

5.9 HCD/HMC: Loading the 2964 processor IOCDS

You now have a production IODF, named `SYS6.IODFB5`. Now you update the IOCDS on the server that you want to upgrade (for example, SCZP401 to SCZP501). The IOCDS will be 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** (Figure 5-117). Ensure that the IODF is the production IODF that was created in 5.8, “HCD: Building the 2964 production IODF” on page 239 and then press Enter.

```
                           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.IODFB5'                          +
```

Figure 5-117 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 5-118).
   Select option **11. Build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes**, and 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
```

Figure 5-118 HCD - Activate or Process Configuration data: build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes

This example assumes that you have connectivity to the 2827 that is being replaced over the HMC local area network to write an IOCDS.

If the server being upgraded is not accessible from the HMC LAN, create an 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 has been written to a USB flash memory drive.

3. The System z Cluster List panel opens (Figure 5-119). In the list, select the 2827 being upgraded by typing a forward slash (/) to update one of its IOCDSs. Then, press Enter.

```
                            System z Cluster List                  Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
/ USIBMSC.SCZP401    2827   H43     SCZP501
# USIBMSC.SCZP501    2964   N63
```

Figure 5-119 HCD - System z Cluster List: selecting processor for IOCDS replace

4. The Actions on selected CPCs panel opens (Figure 5-120). 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)
```

Figure 5-120   HCD - Actions on selected CPCs: work with IOCDSs

5. The IOCDS List panel opens (Figure 5-121). Select the IOCDS that you want to update for the 2827 replacement by typing a forward slash (/) next to it, and 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.SCZP501  IODF35   LPAR   Alternate   No        No          No
_ A1.SCZP501  IODF40   LPAR   POR         No        No          Yes-POR
_ A2.SCZP501  IODF38   LPAR   Alternate   No        No          No
_ A3.SCZP501  IODF39   LPAR   Alternate   No        No          No
```

Figure 5-121   HCD - IOCDS List: selecting IOCDS for replace

6. The Actions on selected IOCDSs panel opens (Figure 5-122). Select option **1. Update IOCDS** and press Enter.

```
--------------- Actions on selected IOCDSs ------


Select by number or action code and press Enter.

1_  1.  Update IOCDS . . . . . . . . . . . . (u)
    2.  Switch IOCDS . . . . . . . . . . . . (s)
    3.  Enable write protection  . . . . . . (e)
    4.  Disable write protection . . . . . . (w)
```

Figure 5-122   HCD - Actions on selected IOCDSs: update IOCDS

7. The Build IOCDSs panel opens (Figure 5-123). 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.IODFB5'

Title1 . IODFB5__________________________________________________________
Title2 : SYS6.IODFB5 - 2014-11-06 17:07

                           Write IOCDS in
IOCDS        Switch IOCDS  preparation of upgrade
A0.SCZP501   No            Yes
***************************** Bottom of data ******************************
```

Figure 5-123 HCD - Build IOCDSs: verifying IODF

Tip: Specifying `Yes` in the Write IOCDS in preparation of upgrade field is required only when you replace the existing hardware and want to write the IOCDS for a 2964 from the existing hardware. The Yes value permits the writing of 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 (Figure 5-124). 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
A0.SCZP501
***************************** Bottom of data ******************************
```

Figure 5-124 HCD - Build IOCDSs: confirm write IOCDS)

9. The Job Statement Information panel is displayed (Figure 5-125). 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 2964 to update its IOCDS.

```
------------------------ Job Statement Information ---------------------

Specify or revise the job statement information.

Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 5-125 HCD - 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 A0 IOCDS REPLACED.

Sev  Msgid    Message Text
 I  CBDA674I  IOCP successfully completed for A0.SCZP501.
```

11. Now if you return to HCD option 2.11 and view the IOCDS, notice that the SNA Address remains at `USIBMSC.SCZP401` (Figure 5-126).

```
                            System z Cluster List                  Row 1 of 4
Command ===> ______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
s USIBMSC.SCZP401    2827   H43     SCZP501
# USIBMSC.SCZP501    2964   N63
****************************** Bottom of data *******************************
```

Figure 5-126 HCD - System z Cluster List: selecting processor for IOCDS verify

12. Also, when you select **USIBMSC.SCZP401**, notice that IOCDS A0 (to which you wrote the upgrade IODF) has a status of `Invalid` (Figure 5-127 on page 246).This error occurs because you specified `Yes` for the Write IOCDS in preparation for upgrade field, and the IOCDS contains IOCP statements and code relevant only for a 2964 processor.

    The status switches when this processor is upgraded to a 2964. The 2964 IOCDS status changes to `Alternate` and the 2827 IOCDSs changes to `Invalid`.

Tip: Generally, rewrite the IOCDS written in preparation for the upgrade at your earliest convenience. Subsequent MESs might cause an IOCDS 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.SCZP501  IODFB5   LPAR   Invalid    No        Yes         No
_ A1.SCZP501  IODF40   LPAR   POR        No        No          Yes-POR
_ A2.SCZP501  IODF38   LPAR   Alternate  No        No          No
_ A3.SCZP501  IODF39   LPAR   Alternate  No        No          No
******************************* Bottom of data ********************************
```

Figure 5-127 HCD - IOCDS List: IOCDS verified with status of Invalid

5.10 HMC: Creating Reset Profile for activation

To activate your Reset Profile using the HMC, complete the steps in this section.

5.10.1 Building the Reset Profile for activation 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 2964 after it is upgraded and handed over from the IBM System Service Representative.

To build the profile, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, or use a remote web browser and select the new 2964.
2. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
3. Under Systems or Members, click the system to select it (in this example, SCZP501).

4. On the Tasks panel, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (Figure 5-128).

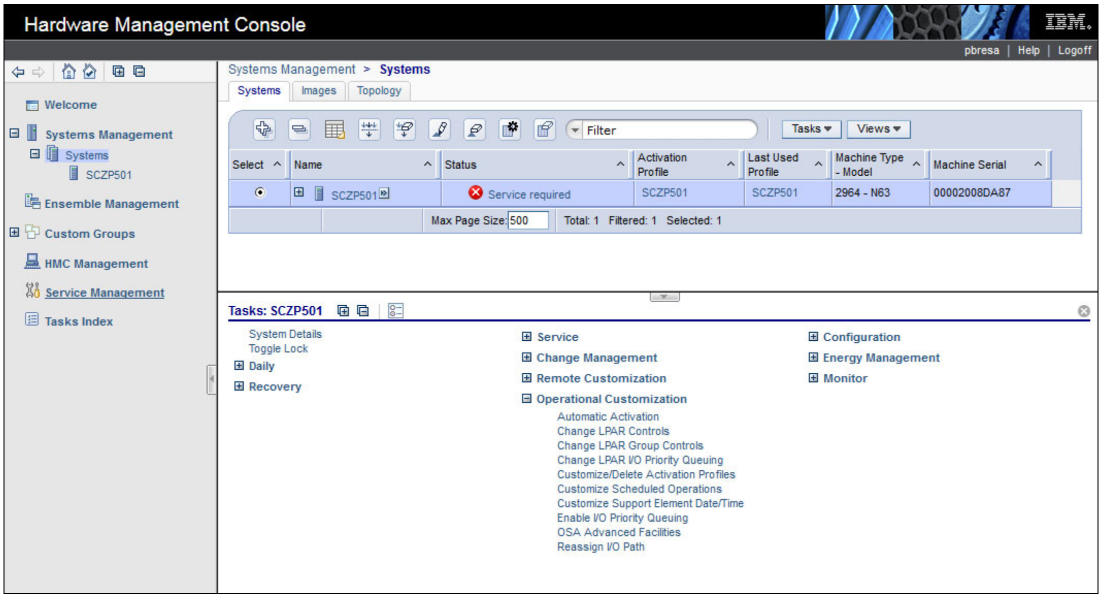

Figure 5-128 HMC - Customize Activation Profiles

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 this new TESTRESET Profile and click **Customize profile**.
8. Click the IOCDS that you updated in the previous step (5.9, "HCD/HMC: Loading the 2964 processor IOCDS" on page 241). The `ACTB0PDL` message is displayed (Figure 5-129).

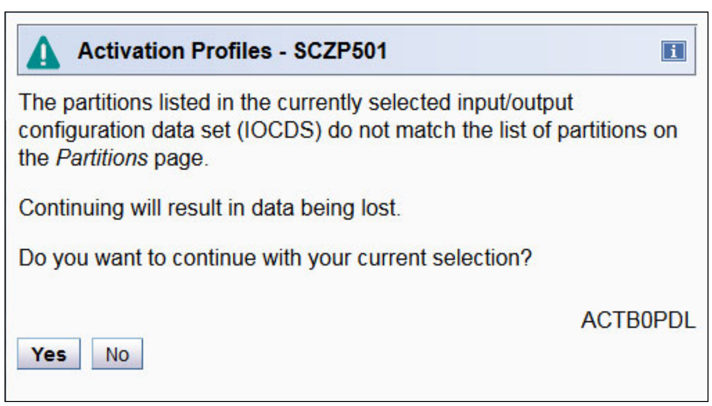

Figure 5-129 HMC - Activation Profiles: ACTB0PDL message

9. Depending on the circumstances, you can answer **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 in the IOCDS that was selected. It also allows you to create new Image profiles for those that it cannot retrieve. In the example, the last option is selected. Click **OK** (Figure 5-130 on page 248).

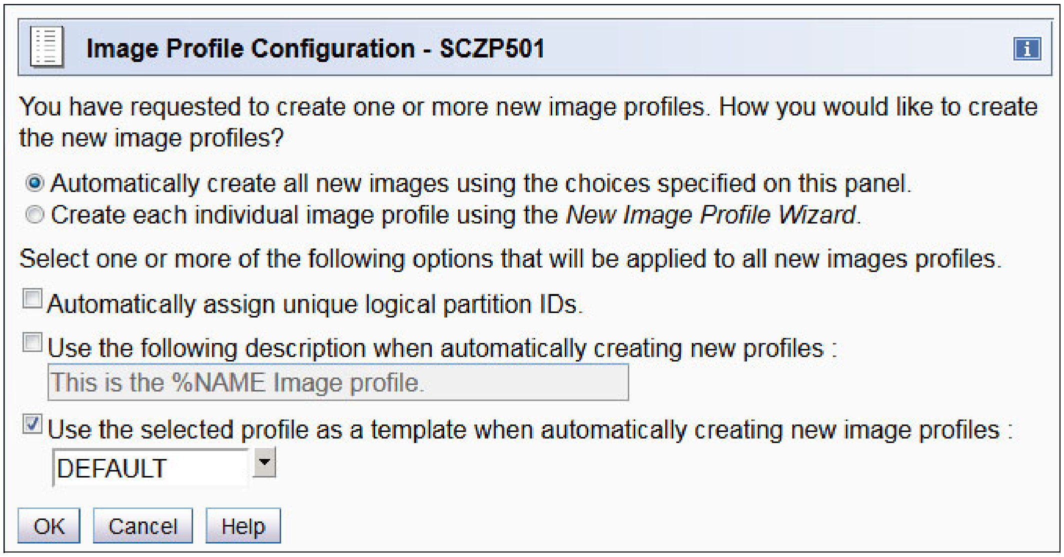

Figure 5-130 HMC - Image Profile automatic build options

11.Note the list of LPARs that were retrieved and built based on the LPARs that were defined in the selected IOCDS. Click **Save** (Figure 5-131).

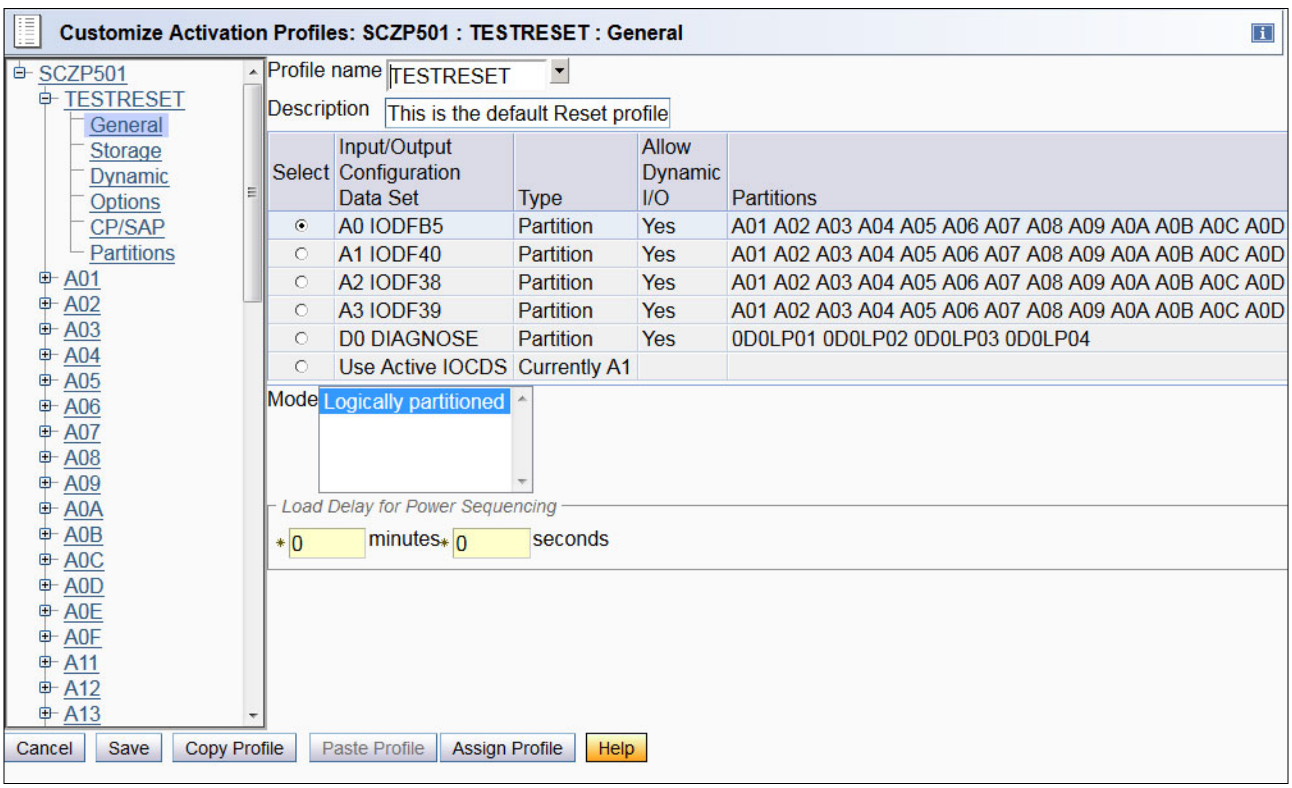

Figure 5-131 HMC - Reset and Image profile list

12.Review the items mentioned in these sections:

– 5.10.2, “Building and verifying image profiles” on page 249
– 5.10.5, “Server Time Protocol (STP) configuration” on page 249

5.10.2 Building and verifying image profiles

While still in the Reset Profile, you can review and update the Image Profile attributes.

5.10.3 Building and verifying Load profiles

Create Load (IPL) profiles by using the DEFAULTLOAD Load profile as a template for all the logical partitions for which you are performing an IPL on this processor.

5.10.4 Building and verifying Loadxx Members in SYS#.IPLPARM

You might need more Loadxx members defined in `SYS#.IPLPARM` for this processor.

If you used the HWNAME parameter to point to the Processor ID, update this parameter to point to the new Processor ID (in this example, from SCZP401 to SCZP501). Sometimes the LPARNAME parameter is also used in the Loadxx members and might need to be reviewed or updated.

5.10.5 Server Time Protocol (STP) configuration

Review the Server Time Protocol (STP) configuration to ensure that the correct External Time Source (ETS) was configured, and that you are connected to the correct Coordinated Time Network (CTN).

For more information about setting up your STP environment, see Chapter 9, “Server Time Protocol (STP)” on page 511.

5.10.6 Running a power-on reset (POR) of the 2964

When the 2817 or 2827 processor is upgraded with a 2964, the IBM System Service Representative will perform 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 in 5.10.1, “Building the Reset Profile for activation and pointing to required IOCDS” on page 246.

The 2964 is now ready to be activated (power-on reset) by using the Production Reset Profile.

6

Replacing an IBM zEC12 with IBM z13

This chapter describes how to replace an existing zEC12 with an IBM z13. Although you can also upgrade from a z196 to an IBM z13, only examples of upgrading a zEC12 are shown in this chapter. An upgrade includes all frames, drawers, and new I/O features.

Because a wide variety of environments exists, the results that you achieve in your environment might differ from those that are described here. Nevertheless, the steps in this chapter provide enough information for you to replicate the approach in your own environment.

This chapter includes the following topics:

- Scenario overview
- HCD: Migrating the existing 2827 IODF
- OSA: Saving and restoring configuration data
- HCD: Validating the 2964 work IODF
- CMT: Assigning PCHIDs to CHPIDs
- HCD: Updating the 2964 work IODF with PCHIDs
- HCD: Building the 2964 production IODF
- HCD/HMC: Loading the 2964 processor IOCDS
- HMC: Steps for profile activation

6.1 Scenario overview

This section describes the replacement scenario.

6.1.1 The configuration process

The ten I/O configuration steps that are described in the “I/O configuration process” topic of *I/O Configuration Using z/OS HCD and HCM,* SG24-7804 are used for the example scenario.

Figure 6-1 shows the general process flow that is followed in this example. The numbered steps are described after the figure.

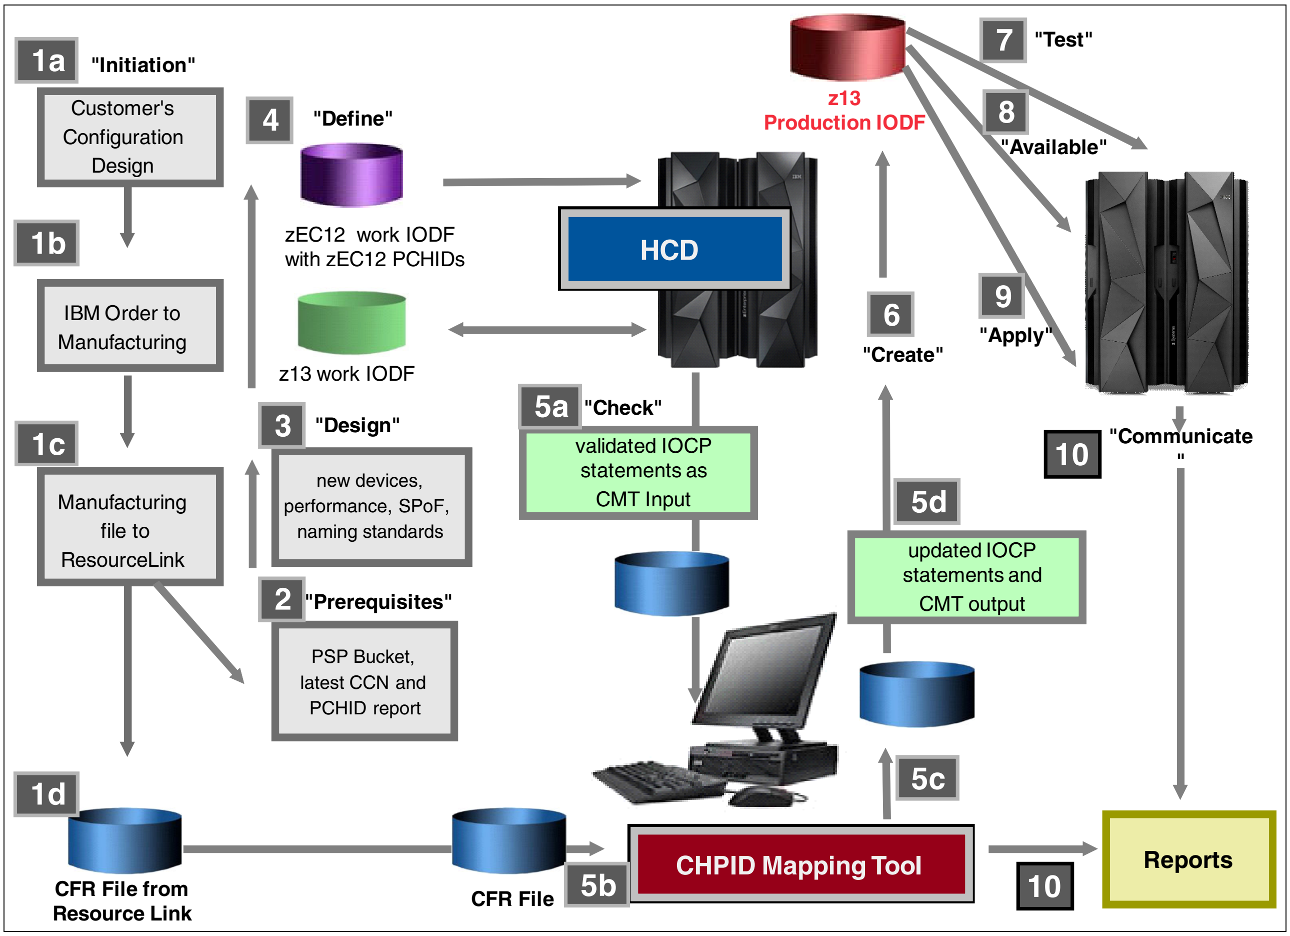

Figure 6-1 Overall configuration process flow

The steps in the figure are as follows:

1. Initiation
   a. When planning to migrate to an IBM z13, the IBM Technical Support team can help you define a configuration design that meets your needs. The configuration is then used during the ordering process.
   b. The IBM order for the configuration is created and passed to the manufacturing process.
   c. The manufacturing process creates a configuration file that is stored at the IBM Resource Link website. This configuration file describes the hardware being ordered. This data is available for download by the client installation team.

   d. A New Order report is created that shows the configuration summary of what is being ordered along with the Customer Control Number (CCN). The CNN can be used to retrieve `CFReport`, which is a data file that contains a listing of hardware configuration and changes for a central processor complex (CPC), from Resource Link.

2. Prerequisites

   Ensure that you have the current PSP Bucket installed. Also, run the SMP/E report with fix category (FIXCAT) exceptions to determine whether any program temporary fixes (PTFs) must be applied. Ensure that you have the most current physical channel ID (PCHID) report and CCN from your IBM System Service Representative.

3. Design

   When you plan your configuration, consider this information:

   – Naming standards
   – FICON switch and port redundancy
   – Adequate I/O paths to your devices for performance
   – OSA channel-path identifier (CHPID) configuration for network and console communications
   – Coupling facility connections internally and to other systems.

   Because the IBM z13 server does not support attachment to the IBM Sysplex Timer, you must consider how the IBM z13 will receive its time source. An IBM z13 cannot join a Coordinated Time Network (CTN) that includes a z10 or earlier as a member. Because the z10 was the last server that supported the IBM Sysplex Timer (9037) connectivity, the z13 cannot be configured as a member of a mixed CTN. The z13 can only join an STP-only CTN.

   When you are planning to replace a z196 or zEC12 with a new z13, plan the replacement of channels that are not supported on IBM z13. You must carefully plan how to replace those, for instance, ISC-3 to HCA3-O or ICA SR for connectivity between z13 and z13.

   You might need to increase CF storage size when you replace z196 or zEC12 with z13. Coupling Facility Control Code (CFCC) level 20 requirements might differ from CFCC level 19 and earlier. Use the CFSizer Tool to get the new CF storage requirements.

4. Define

   The existing z196 or zEC12 I/O configuration is used as a starting point for using Hardware Configuration Definition (HCD). The z196 or zEC12 production input/output definition file (IODF) is used as input to HCD to create a work IODF that becomes the base of the new IBM z13 configuration.

   When the new z13 configuration is added and the obsolete hardware is deleted, a validated version of the configuration is saved in an IBM z13 validated work IODF.

5. Check

   a. From the validated work IODF, create a file that contains the IBM z13 IOCP statements. This IOCP statements file is transferred to the workstation used for the CHPID Mapping Tool (CMT). Hardware Configuration Manager (HCM) can also be used here to transfer the IOCP deck to and from the CMT.

   b. The configuration file that is created by the IBM Manufacturing process in step 1d is downloaded from Resource Link to the CMT workstation.

      The CHPID Mapping Tool (CMT) uses the input data from the files to map logical channels to physical ones on the new IBM z13 hardware.

You might have to make decisions in response to the following situations, among others:

i. Resolving situations in which the limitations on the purchased hardware cause a single point of failure (SPoF). You might must purchase more hardware to resolve some SPoF situations.

ii. Prioritizing certain hardware items over others.

c. After the CMT processing finishes, the IOCP statements contain the physical channels to logical channels assignment that is based on the actual purchased hardware configuration.

The CHPID Mapping Tool (CMT) also creates configuration reports to be used by the IBM System Service Representative and the installation team.

The file that contains the updated IOCP statements created by the CMT, which now contains the physical channels assignment, is transferred to the host system.

d. Use HCD, the validated work IODF file created in step 5a, and the IOCP statements updated by the CMT to apply the physical channel assignments created by the CMT to the configuration data in the work IODF.

6. Create

After the physical channel data is migrated into the work IODF, an IBM z13 production IODF is created and the final IOCP statements can be generated. The installation team uses the configuration data from the IBM z13 production IODF when the final power-on reset is done, yielding an IBM z13 with an I/O configuration ready to be used.

7. Test

IODFs that are modifying existing configurations can be tested in most cases to verify that the IODF is making the intended changes.

8. Available

a. If you are upgrading an existing z196 or zEC12, you might be able to use HCD to write an IOCDS to your system in preparation for the upgrade. If you can write an IOCDS to your current system in preparation for upgrade, do so and let the IBM System Service Representative know which IOCDS to use.

> **Tip:** Using the HCD option “Write IOCDS” in preparation of an upgrade is the preferred method for writing the initial IOCDS when upgrading from a z196 or a zEC12 to an IBM z13. This scenario uses the HCD option *Write IOCDS process*.

b. If the z196 or zEC12 is not network connected to the CPC where HCD is running, or if you are not upgrading or cannot write an IOCDS in preparation for the upgrade, use HCD to produce an IOCP input file. Download this input file to a USB flash drive.

9. Apply

The new production IODF can be applied to the IBM z13 in these ways:

– Using the Power-on Reset process
– Using the Dynamic IODF Activate process

10. Communicate

Communicating new and changed configurations to operations and the appropriate users and departments is important.

6.1.2 Migration path considerations

The migration path from a z196 or zEC12 to an IBM z13 can be either in the form of a field upgrade to the existing z196 or zEC12, or a replacement (push/pull) of an existing z196 or zEC12 with a new IBM z13. Note the following points:

1. A *field upgrade* means that the existing z196 or zEC12 processor serial number is retained during the upgrade.
2. A *replacement* of the existing z196 or zEC12 by a new IBM z13 implies physically removing (*push*) the z196 or zEC12 and bringing in a new IBM z13 (*pull*) to take its place. The replacement IBM z13 has a new serial number that is different from that of the existing z196 or zEC12.

This chapter describes a *replacement* (*push/pull)* scenario.

6.1.3 Planning considerations

The following I/O features can be ordered for a new IBM z13:

- FICON Express16S LX (long wavelength - 10 km) - FC 0418
- FICON Express16S SX (short wavelength) - FC 0419
- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Crypto Express5S - FC 0890
- Flash Express - FC 0403
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171
- PCIe-O SR Integrated Coupling Adapter - FC 0172

The following features, if present in a z196 or zEC12, can be carried forward when you upgrade to an IBM z13:

- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- FICON Express8 LX (long wavelength - 10 km) - FC 3325
- FICON Express8 SX (short wavelength) - FC 3326
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- OSA-Express4S GbE LX (long wavelength) - FC 0404
- OSA-Express4S GbE SX (short wavelength) - FC 0405
- OSA-Express4S 10 GbE LR (long reach) - FC 0406
- OSA-Express4S 10 GbE SR (short reach) - FC 0407
- OSA-Express4S 1000BASE-T Ethernet - FC 0408
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Flash Express - FC 0402

- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171

The following features are not supported on IBM z13:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- OSA-Express2 10 GbE Long Reach
- OSA-Express
- Crypto Express3
- Crypto Express2
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode
- ICB-2
- ICB-3
- ICB-4
- PCIXCC and PCICA
- Parallel channels (use FICON and ESCON converters)

Table 6-1 lists the channel types as described in an input/output configuration data set (IOCDS) that are supported by the z196 (2817), zEC12 (2827), and z13 (2964).

Table 6-1 Channels, links, and adapters with CHPID type and support

Channels	CHPID type	2817 support	2827 support	2964 support
ESCON channels: Connection Channel (ESCON architecture) Channel to Channel (connects to CNC) Converter Channel Path (for BL types) Converter Channel Path (for BY types)	CNC CTC CVC CBY	Up to 240 Up to 240 Up to 240 Up to 240	- No - -	- No - -
FICON native channels that attach to FICON directors or directly to FICON control units: FICON Express 4 SX and LX FICON Express 8 SX and LX FICON Express 8S SX and LX FICON Express 16S SX and LX	FC FC FC FC	Carry forward up to 288 Yes No	Carry forward Carry forward up to 320 No	No up to 64 Carry Forward up to 320
FICON channels that attach to Fibre Channel devices, switches, directors, or Fibre-Channel-to-SCSI bridges	FCP	Yes	Yes	Yes
ISC-3 peer mode channels (connects to another CFP)	CFP	Up to 48	Carry forward	No
IC peer channels (connects to another ICP)	ICP	Up to 32	Up to 32	Up to 32
HCA2-O 12x InfiniBand host channel adapters **HCA2-O 1x LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 32	Up to 32 links Up to 32 links	No No

Channels	CHPID type	2817 support	2827 support	2964 support
HCA3-O 12x InfiniBand host channel adapters **HCA3-O LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 48	Up to 32 links Up to 64 links	Up to 32 links Up to 64 links
PCIe-O SR Integrated Coupling Adapter	CS5	No	No	Up to 32 links
HiperSockets (IQDIO) channels	IQD	Up to 32	Up to 32	Up to 32
OSA- Express2 GbE LX/SX	OSD and OSN	Up to 48 ports carried forward	No	No
OSA- Express2 1000BASE-T	OSE, OSD, OSC, and OSN	Yes	No	No
OSA- Express2 10 GbE LR	OSD	Yes	No	No
OSA-Express3 GbE LX/SX	OSD and OSN	Up to 96 ports	Carry forward	No
OSA-Express3 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	OSC, OSD, OSE, and OSN: Yes OSM: No	Carry forward	No
OSA- Express3 10 GbE LR/SR	OSD and OSX	OSD: Yes OSX: No	Carry forward	No
OSA-Express4S GbE LX/SX	OSD	Up to 96 ports	Up to 96 ports	Carry forward
OSA-Express4S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Carry forward
OSA-Express4S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Carry forward
OSA-Express5S GbE LX/SX	OSD	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Up to 48 ports
10GbE RoCE Express SR	ROCE	No	Up to 32 ports	Up to 32 ports
zEDC Express Enterprise Data Compression	ZEDC	No	Max 8 features	Max 8 features

Consider the following information when you plan your configuration:

- Coupling links
- Hardware Management Console (HMC)
- Software support
- Open Systems Adapter - Integrated Console Controller
- Fibre Channel Protocol
- CPC name versus Processor ID
- Local system name

Coupling links

Only the following coupling facility CHPIDs are supported:

- CHPID type CIB: PSIFB links connecting to an HCA3-O 12x and HCA3-O LR 1x
- CHPID type ICP: Internal Coupling links
- CHPID type CS5: Integrated Coupling adapter PCIe-O SR

Considerations: Coupling links can be defined as both Coupling and STP links, or STP-only links. IBM z13 does not support the ISC-3 features.

Hardware Management Console (HMC)

The HMC can appear either as the current HMC does, or as an HMC that can run code to manage an Ensemble. The current HMC is used to manage, monitor, and operate one or more IBM z Systems servers and their associated logical partitions. An HMC that has ensemble code running is an HMC attached to one or more zEnterprise Systems configured as ensemble members. A particular ensemble is managed by a pair of HMCs in primary and alternate roles.

The HMC has a global (Ensemble) management function, whereas the SE has local node management responsibility. When tasks are performed on the HMC, the commands are sent to one or more Support Elements (SEs), which then issue commands to their CPCs.

The IBM z13 requires HMC Application V2.13.0 (driver level 21) or later, and uses only Ethernet for its network connection. The HMC and the SEs do not contain a floppy disk drive. Therefore, you must use a USB flash memory drive to input and back up client configuration data.

Software support

HCD V2.1 and later with the Preventive Service Planning (PSP) bucket for 2964DEVICE and PTFs) is required to define and support some of the new features of the z13.

Open Systems Adapter - Integrated Console Controller

Because support has been withdrawn for the 2074 console controllers, consider using OSA-Express4S 1000BASE-T or OSA-Express5S1000BASE-T CHPIDs defined as TYPE=OSC. With this OSA card, you can set up console function that is supported by a configuration file defined on the Support Element for that processor.

Fibre Channel Protocol

When you use CHPIDs defined as TYPE=FCP, consider N-Port ID Virtualization (NPIV).

For more information about FCP CHPIDs and the WWPN prediction tool to manage them, see “Batch Network Analyzer and Compression” on page 47.

CPC name versus Processor ID

HCD allows you to define different processors (logical) to the same CPC (physical). The Processor ID must be unique within the same IODF, but the CPC name does not. Therefore, the CPC name does not need to match the Processor ID. This advantage is useful when you have several processor/logical partition/control unit setups that share a physical CPC within the same IODF. Furthermore, the Processor ID is what is defined for the optional HWNAME parameter in the LOAD member in SYS1.IPLPARM.

The CPC name is coded in HCD option 1.3 under View Processor Definition in the CPC name field under SNA address, along with a Network name. It is the CPC name, and not the Processor ID, that is displayed on the HMC.

When you view the network information for a CPC through the HMC, the SNA address is made up of a network name and CPC name separated by a dot (for example, `USIBMSC.SCZP501`). These values are defined in the Support Element for the CPC. They must match the values that are set in the IODF so that HCD option 2.11 can find the CPC to write an IOCDS in the System z Cluster List.

Local system name

An extra system name, LSYSTEM, is used to identify the local system name of a server when you define PSIFB type=CIB or ICA type=CS5 coupling links.

This data field can be found when you change a CIB-capable processor under HCD option 1.3.

The LSYSTEM field can be set or changed to any one to eight alphanumeric characters. Also, it can begin with either an alphabetic or numeric character. All characters are uppercase.

The following rules determine whether, and where, HCD sets the LSYSTEM keyword automatically:

- If a CIB-capable processor is defined and the CPC name is set but the local system name is not set, HCD sets the local system name to the CPC name.
- If a CIB-capable processor that has not yet defined a CPC name is changed to obtain a CPC name but no local system name, HCD sets the CPC name to the local system name.
- If a non-CIB capable processor is changed to a support level that is CIB capable, and the processor has a CPC name set but no local system name, the local system name defaults to the CPC name.
- If the processor definition is changed such that the local system name is explicitly removed, HCD does not do any defaulting.
- If a processor has a local system name set (whether it has a CPC name or not), any change to the local system name must be done explicitly. There is no implicit name if the CPC name or the support level is changed.
- During Build Production IODF, verification determines whether a local system name is set for the processor if the processor has a CIB channel path defined. If this verification fails, an error message is displayed and the production IODF is not built.

Generally, set the local system name the same as the CPC name.

The following additional keywords are for the ID statement in an IOCP deck:

CPATH Specifies the CSSID and CHPID on the connecting system.

CSYSTEM Connecting System name (LSYSTEM in processor definition).

AID Host Card Adapter ID (found in PCHID report).

PORT Host Card Adapter port.

6.1.4 Replacement (push/pull) scenario

This scenario describes the configuration steps to replace an existing 2827 processor with a new 2964 processor. Key factors are as follows:

- HCD requires a new Processor ID for the 2964.
- HCD requires a new CPC name for the 2964.
- The 2964 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 2827 production IODF.
- The target IODF is a new 2964 work IODF.
- HCD actions are as follows:
  - Migrate updated IOCP statements.
  - Build production IODF.
  - Remote write IODF to IOCDS.
- HMC actions are as follows:
  - Build Reset Profile and point to required IOCDS.
  - Build/Verify Image Profiles.
  - Build/Verify Load Profiles.
  - Run a power-on reset.

The example uses a 2827-H43 with a Processor ID of SCZP401 and four CSSs (CSS ID=0, CSS ID=1, CSS=2, and CSS ID=3). This system is replaced with a 2964-N63 with a Processor ID of SCZP501 and six CSSs.

The CPC name SCZP401 and serial number 02-0B8D7 are changed to the CPC name of SCZP501 and serial number of 02-8DA87.

The following CHPID types are migrated:

- OSD, OSC, OSM, and OSX
- FC and FCP
- CIB and ICP
- IQD

The following hardware or CHPID types are *not* supported and *not* migrated to the 2964:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- Crypto Express3
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode

Table 6-2 summarizes the migration options and tool requirements. The process steps are described in 6.2, “HCD: Migrating the existing 2827 IODF” on page 261.

Table 6-2 2827 I/O configuration migrated to a 2964

2827 to 2964	Replace an existing 2827 with a 2964 (push/pull)
Processor ID	Required to change the processor ID to a new ID
CPC name	New CPC name
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 (see step 2 on page 261)
CMT Program (needed or not)	Optional, but generally a good idea
CFReport file (CCN) (needed or not)	Required
IOCP (import from validated work IODF)	Yes
CMT Actions (PCHID reset)	Yes
CMT IOCP Output	Yes
CMT Reports	Yes (CHPID and CHPID to CU Report)

6.2 HCD: Migrating the existing 2827 IODF

The following steps explain how to define the existing 2827 I/O configuration to the new 2964 IODF by using HCD. Then migrate the channel subsystem and logical partitions from the 2827 to the 2964 server. Using HCD, the sequence of operations is as follows:

1. Create a work IODF from the current 2827 production IODF.
2. Repeat the 2827 processor that is being replaced.
3. Observe the Coupling Link messages for later reference.
4. Delete unsupported items from the repeated 2827.
5. Change the repeated 2827 to a 2964, and then delete the 2827.
6. Redefine all required CF connections to other processors and any Internal CF connections required.
7. Define any additional CHPIDs, control units, and devices you might be adding in the replacement.
8. When necessary, overdefine channel paths.
9. Save and restore the OSA configuration information.
10. Build a validated 2964 work IODF.
11. Create an IOCP statements file and transfer it to the workstation that is running the CHPID Mapping Tool. This step can also be performed with HCM.

12. Import CFReport and IOCP statements into the CMT.
13. Perform hardware resolution and PCHID/CHPID availability.
14. Create configuration reports for yourself and your IBM System Service Representative.
15. Import IOCP statements that have been updated with PCHIDs back into the validated work IODF.
16. Build the production IODF.
17. Remotely write the IOCP to an IOCDS on the 2964 Support Element. If this action is not possible, copy the IOCP statements to a USB memory flash drive and run the Stand-Alone Input/Output Configuration Program to load the IOCP statements to an IOCDS on the 2964 Support Element.
18. Build the Reset, Image, and Load Profiles, if required.
19. Perform a power-on reset (activate) of 2964.

6.2.1 Creating the work IODF from the current 2827 production IODF

Create a work IODF from an existing production IODF. To do this use the HCD primary panel and enter the name of the current production IODF that contains the 2827 processor that you are replacing in the I/O definition file entry field (for example, `SYS6.IODFB3`).

6.2.2 Repeating the 2827 processor that is being replaced

To repeat the 2827 processor that is being replaced, complete the following steps:

1. From the main HCD panel, select option **1.3. Processor List**.
2. In the Processor List (Figure 6-2), enter **r** (for repeat) next to the 2827 that you want to replace and press Enter.

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
r SCZP401  2827     H43      LPAR  00B8D72827 Helix
****************************** Bottom of data ******************************
```

Figure 6-2 HCD - Processor List: repeating a processor

3. The Identify Target IODF panel is displayed. Complete one of the following actions:
   - To retain all the other processor definitions in the IODF, press Enter.
   - Enter a Target IODF data set name. In this case, only the processor 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.IODFB6.WORK`).
5. The Repeat Processor panel is displayed(Figure 6-3 on page 263). Enter the Processor ID of the new 2964 (in this example, `SCZP501`), keep all the other fields unchanged, and press Enter to continue.

```
+---------------------------- Repeat Processor -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . . . . . . SCZP501_                                   |
|                                                                           |
| Processor type . . . . . . . : 2827                                       |
| Processor model  . . . . . . : H43                                        |
| Configuration mode . . . . . : LPAR                                       |
|                                                                           |
| Serial number  . . . . . . . . 00B8D72827                                 |
| Description  . . . . . . . . . Helix                                      |
|                                                                           |
| Specify SNA address only if part of an System z cluster:                  |
|                                                                           |
| Network name . . . . . . . . . USIBMSC   +                                |
| CPC name . . . . . . . . . . . SCZP401   +                                |
|                                                                           |
| Local system name  . . . . . . SCZP401                                    |
|                                                                           |
+-----------------------------------+ F4=Prompt   F5=Reset    F9=Swap       |
| New IODF SYS6.IODFB6.WORK defined. |                                      |
+-----------------------------------+ --------------------------------------+
```

Figure 6-3 HCD - Repeat Processor: defining new processor ID

6.2.3 Coupling Link information messages

You might receive severity messages (E, I, or W). As shown in Figure 6-4, CBDG441I (severity I) messages are received because the coupling link CHPIDs are not being copied to the 2964 definition.

```
------------------------------ Message List ------------------------------
  Save  Query  Help
 --------------------------------------------------------------------------
                                                               Row 1 of 242
 Command ===> ___________________________________________ Scroll ===> CSR

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

 / Sev Msg. ID  Message Text
 _ I   CBDG441I The coupling facility connection between channel path
 #              0.80 of processor SCZP401 and channel path 0.90 of
 #              processor SCZP201 is not copied.
 _ I   CBDG441I The coupling facility connection between channel path
 #              0.81 of processor SCZP401 and channel path 1.A0 of
 #              processor SCZP401 is not copied.
 _ I   CBDG441I The coupling facility connection between channel path
 #              0.82 of processor SCZP401 and channel path 1.A2 of
 #              processor SCZP401 is not copied.
 _ I   CBDG441I The coupling facility connection between channel path
 #              0.84 of processor SCZP401 and channel path 0.98 of
```

Figure 6-4 HCD - Message List: showing CBDG441I

To resolve this problem, complete the following steps:

1. Scroll until you reach the end of the messages and see the `CBDG271I requested action on object SCZP501 successfully processed` message.
2. Press F3 or F12 to continue. As shown in Figure 6-5, there is now a new 2964 processor named `SCZP501` on `SYS6.IODFB6.WORK`.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2827     H43      LPAR  00B8D72827 Helix
******************************* Bottom of data *******************************
```

Figure 6-5 HCD - Processor List: with repeated processor

6.2.4 Changing the 2827 to a 2964 and deleting the 2827

You can either keep the original copy of the 2827 (SCZP401) or delete it from the IODF. In this example, keep it in the IODF for now.

To change the 2827 to a 2964, complete the following steps:

1. Enter `c` (change) next to SCZP501 to change the 2827 to a 2964 and press Enter. The Change Processor Definition panel opens (Figure 6-6 on page 265).
2. Make the following updates and press Enter:
   - Update `Processor type` to `2964`.
   - Update `Processor Model` to `N63`.
   - Update `Serial Number` to `08DA872964`.
   - Update the Description if you desire
   - Update `CPC name` to `SCZP501`.
   - Update `Local system name` to `SCZP501`.

```
---------------------- Change Processor Definition ---------------------


Specify or revise the following values.

Processor ID . . . . . . . . : SCZP501
Support level:
XMP, 2827 GA2 support
Processor type . . . . . . . . 2964      +
Processor model  . . . . . . . N63       +
Configuration mode . . . . . . LPAR      +

Serial number  . . . . . . . . 08DA872964  +
Description  . . . . . . . . . Sphinx

Specify SNA address only if part of an System z cluster:

Network name . . . . . . . . . USIBMSC   +
CPC name . . . . . . . . . . . SCZP501   +

Local system name  . . . . . . SCZP501
```

Figure 6-6 HCD - Change Processor Definition: changing repeated processor values

3. Press Enter. The example shows PCHID numbers exceed the maximum value for a Processor type of 2964. Because we will be remapping the PCHIDs to CHPIDs in CHPID Mapping Tool section of this chapter, we can either manually edit the PCHIDs for this Processor definition to fit into the allowed PCHID range, or manually delete them in the IODF (Figure 6-7).

```
----------------------------- Message List ------------------------------
  Save  Query  Help
--------------------------------------------------------------------------
                                                                Row 1 of 165
Command ===> ___________________________________________ Scroll ===> CSR

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

/ Sev Msg. ID  Message Text
_ E   CBDG522I The maximum value 4FF for the physical channel ID has
#              been exceeded for CHPID / function 0.00 of processor
#              SCZP501. Actual value: 574
_ E   CBDG522I The maximum value 4FF for the physical channel ID has
#              been exceeded for CHPID / function 0.01 of processor
#              SCZP501. Actual value: 54C
_ E   CBDG522I The maximum value 4FF for the physical channel ID has
#              been exceeded for CHPID / function 0.02 of processor
#              SCZP501. Actual value: 58C
_ E   CBDG522I The maximum value 4FF for the physical channel ID has
#              been exceeded for CHPID / function 0.03 of processor
```

Figure 6-7 HCD - Message List: showing CBDG22I

4. Press PF12 to return to the processor list. In this example we will be manually removing the PCHID definitions for the repeated processor SCZP501. Once this has been completed, perform the change processor type step again.

5. Press Enter. The Update Channel Path Identifiers panel is displayed (Figure 6-8). In the example, no changes are made.

```
-------------------- Update Channel Path Identifiers ---------------------
                                                                Row 1 of 88
Command ===> __________________________________________ Scroll ===> CSR

Specify any changes to the channel path identifiers in the list below.

Processor ID . . . . : SCZP501      Sphinx
Channel Subsystem ID : 0

CHPID  Type  Side  Until CHPID   New CHPID +
00     OSD         __            00
01     OSC         __            01
02     OSD         __            02
03     OSD         __            03
04     OSD         __            04
06     OSD         __            06
07     OSD         __            07
0C     OSD         __            0C
0D     OSC         __            0D
12     OSD         __            12
13     OSD         __            13
```

Figure 6-8 HCD - Update Channel Path Identifiers: not changed

6. Press Enter *for each of the Channel Subsystem IDs*.
7. HCD displays the following messages because the 2964 has more configuration resources than the 2827 processor definition. Press PF12 to return to the processor list.

```
------------------------------ Message List -------------------------------
  Save  Query  Help
--------------------------------------------------------------------------
                                                               Row 1 of 4
Command ===> ___________________________________________ Scroll ===> CSR

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

/ Sev Msg. ID  Message Text
_ I   CBDG572I Virtual functions for function type ROCE have been
#              changed.
_ I   CBDG540I Definition of processor SCZP501 has been extended to its
#              maximum configuration.
**************************** Bottom of data *****************************
```

Figure 6-9 HCD - Processor List: maximum configuration extended

The repeated 2827 processor is successfully changed to a 2964-N63 (Figure 6-10).

```
 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
****************************** Bottom of data *******************************
```

Figure 6-10 HCD - Processor List: changed processor

6.2.5 Deleting the 2827 processor definition

Now that the 2827 is repeated and changed to a 2964, the original 2827 definition (SCZP401) must now be deleted so that the required Coupling Links can be restored.

To delete the processor definition, complete the following steps:

1. Enter `d` (for delete) next to the SCZP401 processor in the Processor List (Figure 6-11).

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
d SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
****************************** Bottom of data *******************************
```

Figure 6-11 HCD - Processor List: deleting processor

2. Press Enter to confirm the deletion of the processor (Figure 6-12).

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
****************************** Bottom of data *******************************
```

Figure 6-12 HCD - Processor List: with processor deleted

6.2.6 Reconnecting the Coupling Link channel paths that were not migrated

Next, manually redefine the Coupling Links you want from the SCZP501 processor to any other processor, along with any Internal Coupling Facility links you want. To help in this effort, you can get a CF connection report from the previous production IODF containing the 2827. Alternatively, you can make a note of all CBDG441I error messages that you received in 6.2.5, “Deleting the 2827 processor definition” on page 267.

6.2.7 Defining more I/O

Define any additional CHPIDs, control units and devices, FICON channel-to-channel connections (FCTCs), and so on, that you might be adding into the 2964 during the replace.

This section shows use of HCD and HCM to define the following CHPIDs and PCIe functions:

- Defining a CF link over CHPID type CS5 using HCD
- Defining a CF link over CHPID type CS5 using HCM
- Defining a RoCE PCIe function using HCD
- Defining a RoCE PCIe function using HCM
- Defining a zEDC EXPRESS PCIe function using HCD
- Defining a zEDC-EXPRESS PCIe function using HCM

Defining a CF link over CHPID type CS5 using HCD

A new CHPID type (CS5) is introduced with the z13 and uses a new PCIe-O SR ICA hardware card for the z13 named *PCIe-O SR Integrated Coupling Adapter FC 0172*.

It runs at 8 GBps and up to a length of 150 meters. Presently its connectivity is supported only between z13 and another z13.

The HCD definition process is similar to CHPIDs of type=CIB except that the type=CS5.

In this example, we define these items:

- CHPID=0.86 to AID/Port ID=27/1 to processor=SCZP501
- CHPID=0.96 to AID/Port ID=27/2 to processor=SCZP502
- A CF link between these two CS5 CHPIDs

Complete these steps:

1. From the main HCD panel, select option **1.3. Processor List**. Enter `s` (work with attached channel paths) next to the first CPC (SCZP501) to which you want to define the CS5 CHPID and press Enter.
2. From the Channel Subsystem List, enter `s` (work with attached channel paths) next to the CSSID that you want. In our example, we use CSSID=0. Press Enter.
3. Enter `add` on the command line in the Channel Path List panel to add a new channel and then press Enter.
4. Enter the following values in the Add Channel Path panel (Figure 6-13 on page 269) and press Enter:
   – Channel path ID: `86`
   – Channel path type: `CS5`
   – Operational mode: `SHR`

```
-------------------------- Add Channel Path --------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 86    +           Channel ID  ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +
Description  . . . . . . CS5 CHPID definition___________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 6-13 HCD - Add Channel Path: defining a CS5 CHPID

5. Enter the following values in the Specify HCA Attributes panel (Figure 6-14) and then press Enter:
   - Adapter ID of the HCA: 27
   - Port on the HCA: 1

```
---------------- Specify HCA Attributes -------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 1   +
```

Figure 6-14 HCD - Specify HCA Attributes: defining CS5 HCA attributes

6. Select the required Access LPARs for CHPID access list. In our example, we use LPAR A0C (CF) and A01, A02, and A03 (OS). Press Enter. See Figure 6-15 on page 270.

```
-------------------------- Define Access List ---------------------------
                                                              Row 1 of 15
Command ===> ________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 86     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      OS
_ 0      A0B              B      CF/OS CHPID holder
/ 0      A0C              C      CF    CF7A
_ 0      A0D              D      CF    CF7B
_ 0      A0E              E      CF    CF7C
_ 0      A0F              F      CF    CF7D
/ 0      A01              1      OS    SC76
/ 0      A02              2      OS    SC74
/ 0      A03              3      OS    SC75
_ 0      A04              4      OS    ZZ0VM1
```

Figure 6-15 HCD - Define Access List: selecting partition for CS5 CHPID access

7. Select any Candidate LPARs for the CHPID access list. In our example we do not select any candidate LPARs. Press Enter.
8. From the main HCD panel, select option **1.3. Processor List**. Enter `s` (work with attached channel paths) next to the second CPC (SCZP502) to which you want to define the CS5 CHPID and press Enter.
9. From the Channel Subsystem List, enter `s` (work with attached channel paths) next to the desired CSSID (in our example, we use CSSID=0), and then press Enter.
10. Enter `add` on the command line in the Channel Path List panel to add a new channel and then press Enter.
11. Enter the following values in the Add Channel Path panel (Figure 6-16 on page 271) and press Enter:
    - Channel path ID: `96`
    - Channel path type: `CS5`
    - Operational mode: `SHR`

```
--------------------------- Add Channel Path ------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP502      test machine definition
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 96    +           Channel ID  ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +
Description  . . . . . . CS5 CHPID definition____________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 6-16 HCD - Add Channel Path: defining a CS5 CHPID

12. Enter the following values in the Specify HCA Attributes panel (Figure 6-17) and then press Enter:
   – Adapter ID of the HCA: 27
   – Port on the HCA: 2

```
----------------- Specify HCA Attributes ------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 2   +
```

Figure 6-17 HCD - Specify HCA Attributes: defining CS5 HCA attributes

13. Select the required Access LPAR for CHPID access list. In our example, we use LPAR A0D (CF) and A01, A02, and A03 (OS). Then press Enter. See Figure 6-18 on page 272.

```
-------------------------- Define Access List ---------------------------
                                                              Row 1 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 96     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      CF/OS
_ 0      A0B              B      CF/OS
_ 0      A0C              C      CF
/ 0      A0D              D      CF
_ 0      A0E              E      CF/OS
_ 0      A0F              F      CF/OS
/ 0      A01              1      OS
/ 0      A02              2      OS
/ 0      A03              3      OS
_ 0      A04              4      CF/OS
```

Figure 6-18 HCD - Define Access List: selecting partition for CS5 CHPID access

14. Select any Candidate LPARs for the CHPID access list. In our example, we do not select any candidate LPARs. Press Enter.
15. Return to the first CHPID definition (0.86 on processor SCZP501) and type **f** (connect CF channel paths) next to CHPID 0.86 and press Enter. See Figure 6-19.

```
                              Channel Path List    Row 91 of 117 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f 86    27/1  CS5   SHR   __  __ __   N   No  CS5 CHPID definition
_ 87    34/1  CS5   SHR   __  __ __   N   No  P Fram
_ 90    08/2  CIB   SHR   __  __ __   Y   No  12x P201 T/O
_ 94    0B/2  CIB   SHR   __  __ __   Y   No  1x
_ 95    0F/2  CIB   SHR   __  __ __   Y   No  1x
```

Figure 6-19 HCD - Channel Path List: connecting a CF channel path

16. Type **p** (CF channel path connectivity list) next to CHPID 0.86 and press Enter. See Figure 6-20.

```
                     CF Channel Path Connectivity List            Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N     SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N     SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N     SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N     SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N     SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N     SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
p 86  27/1  Y  CS5  SHR  N
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N     SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 6-20 HCD - CF Channel Path Connectivity List: creating a CF link

17. Enter the following values in the Connect to CF Channel Path panel (Figure 6-21) and then press Enter:

– Destination processor ID: `SCZP502`
– Destination channel subsystem ID: `0`
– Destination channel path ID: `96`

```
--------------- Connect to CF Channel Path --------


Specify the following values.

Source processor ID  . . . . . : SCZP501
Source channel subsystem ID  . : 0
Source channel path ID . . . . : 86
Source channel path type . . . : CS5

Destination processor ID . . . . . . SCZP502_  +
Destination channel subsystem ID . . 0   +
Destination channel path ID  . . . . 96  +

Timing-only link . . . . . . . . . . No
```

Figure 6-21 HCD - Connect to CF Channel Path: specifying the destination attributes

18. Accept or override the control unit and device numbers for processor *SCZP501*, then press Enter. See Figure 6-22.

```
---------------- 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 . . . . . . . : SCZP501
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 86         Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF6  +

Device number  . . . . . . . FE3D
Number of devices  . . . . . 8
```

Figure 6-22   HCD - Add CF Control Unit and Devices: accept or override default values

19. Accept or override the control unit and device numbers for processor *SCZP502*, then press Enter. See Figure 6-23.

```
---------------- 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 . . . . . . . : SCZP502
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 96         Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF5  +

Device number  . . . . . . . FE2D
Number of devices  . . . . . 8
```

Figure 6-23   HCD - Add CF Control Unit and Devices: accept or override default values

HCD returns to the CF Channel Path Connectivity List panel (Figure 6-24) where you can see the CF link now defined.

```
                     CF Channel Path Connectivity List             Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N     SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N     SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N     SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N     SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N     SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N     SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
_ 86  27/1  Y  CS5  SHR  N     SCZP502.0  96  27/2  Y  CS5  SHR    CFP   8
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N     SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 6-24 HCD - CF Channel Path Connectivity List: CF link now created

Defining a CF link over CHPID type CS5 using HCM

A CHPID type (CS5) was introduced with the IBM z13 and uses a new PCIe-O SR ICA hardware card for the IBM z13 named *FC 0172 PCIe-O SR Integrated Coupling Adapter*.

It runs at 8 GBps and up to a length of 150 meters. Presently, its connectivity is supported only between one IBM z13 and another IBM z13.

The HCM definition process is similar to CHPIDs of type=CIB except that the type=CS5.

In this example, we define these items:

- CHPID=0.86 to AID/Port ID=27/1 to processor=SCZP501
- CHPID=0.96 to AID/Port ID=27/2 to processor=SCZP502
- A CF link between these two CS5 CHPIDs

Complete the following steps:

1. From the HCM main page, click **Edit** → **Processor.**
2. Select the first CPC (SCZP501), and then click **OK**.
3. Click once on the desired CSS (CSSID=0), and then click **CHPIDs.**
4. Click **Create** HCM. The Create CHPIDs window opens.
5. Enter the following values in the Create CHPIDs window, and then click **OK** (Figure 6-25 on page 276).
   - Channel path ID: `86`
   - Channel path type: `CS5`
   - Operational mode: `SHR`

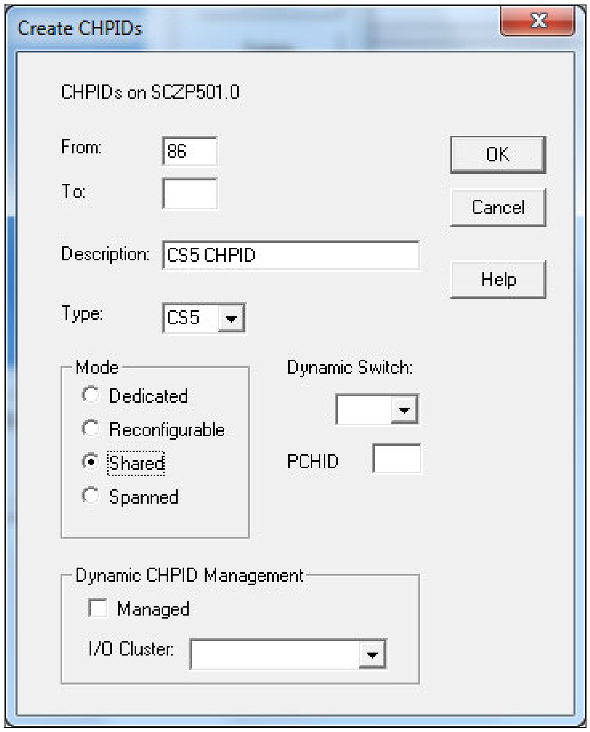

Figure 6-25 HCM - Create CHPIDs: defining a CS5 CHPID

6. Enter the following values in the Host Communication Adapter Attributes window, and then Click **OK** (Figure 6-26):

   – Adapter ID of the HCA: `27`
   – Port on the HCA: `1`

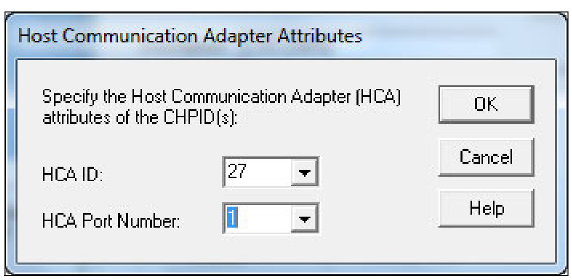

Figure 6-26 HCM - Specify HCA Attributes: defining CS5 HCA attributes

7. Select the required Access and Candidate LPARs for CHPID access list. In our example, we use LPAR A0C (CF) and A01, A02, and A03 (OS), and no Candidate LPARs. Then click **OK** (Figure 6-27).

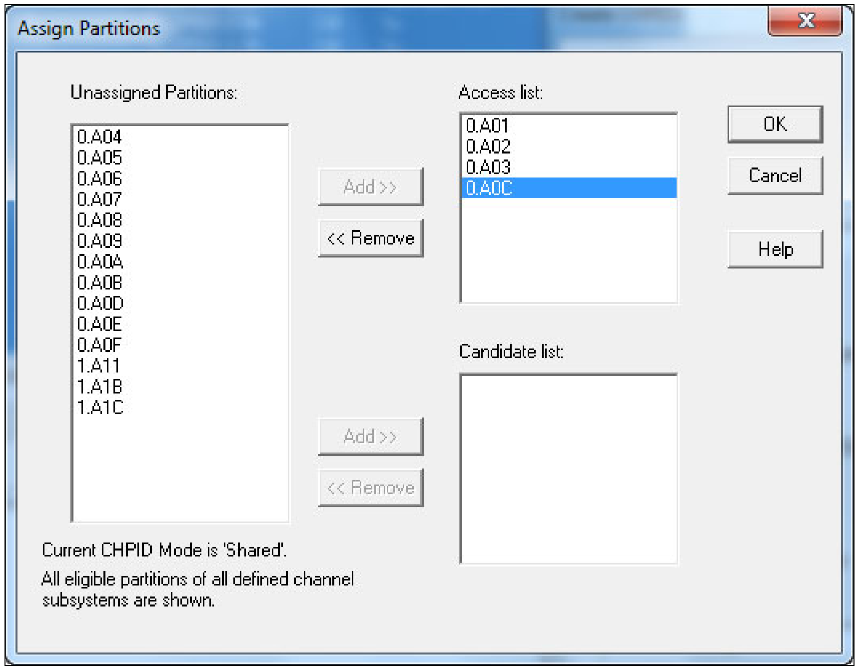

Figure 6-27 HCM - Assign Partitions: selecting partition for CS5 CHPID access

8. Now observe the CS5 CHPID defined in the CHPID summary list (Figure 6-28).

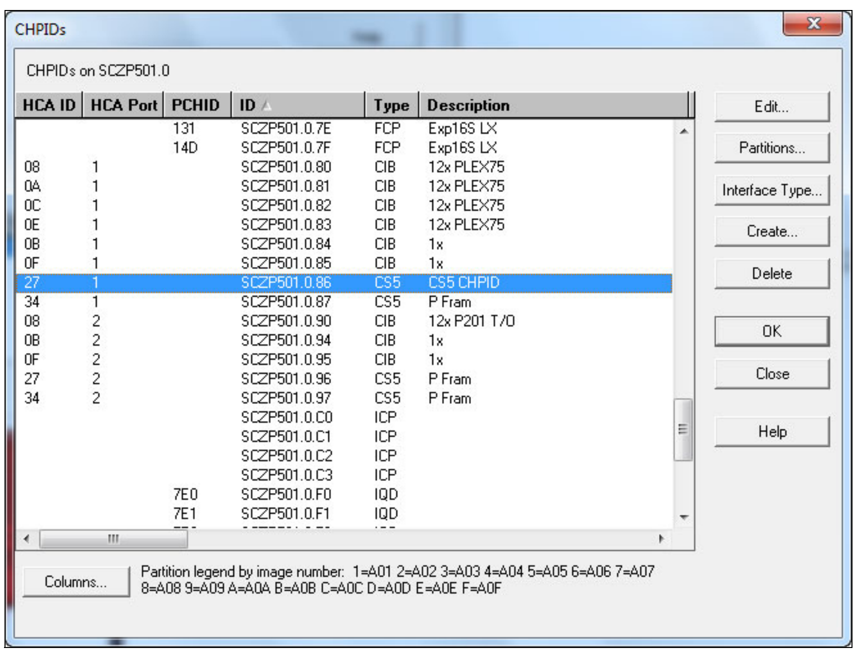

Figure 6-28 HCM - CHPIDs: first CS5 CHPID defined

9. From the HCM main page, click **Edit** → **Processor**.

10. Select the second CPC (SCZP502), and then click **OK**.

11. Click once on the desired CSS (CSSID=0), and then click **CHPIDs**.
12. Click **Create** HCM. The Create CHPIDs window opens.
13. Enter the following values in the Create CHPIDs window, and then click **OK**:
    - Channel path ID: `96`
    - Channel path type: `CS5`
    - Operational mode: `SHR`
14. Enter the following values in the Host Communication Adapter Attributes window, and then click **OK**:
    - Adapter ID of the HCA: `27`
    - Port on the HCA: `2`
15. Select the required Access and Candidate LPARs for CHPID access list. In our example we use LPAR A0D (CF) and A01, A02 and A03 (OS), and no Candidate LPARs. Then click **OK**.
16. Now observe the CS5 CHPID defined in the CHPID summary list (Figure 6-29).

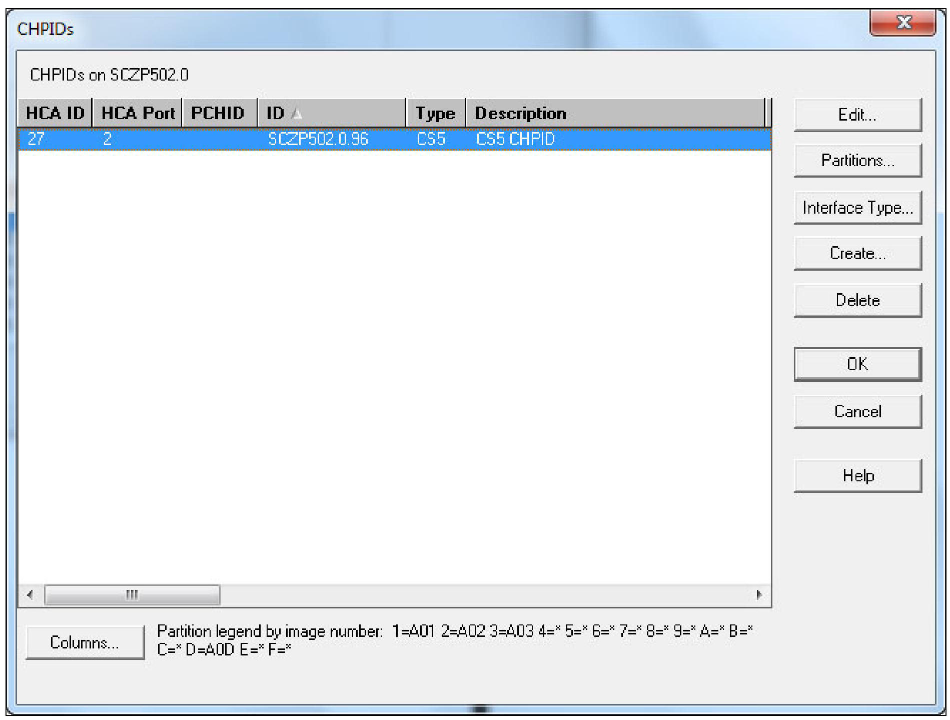

Figure 6-29 HCM - CHPIDs: second CS5 CHPID defined

17. From the HCM main page, click **Create** → **CF connection.**

18.HCM displays the Create Coupling Facility Connection window (Figure 6-30). From the pull-down, select the first CS5 CHPID that we defined (SCZP501 - 0.86).

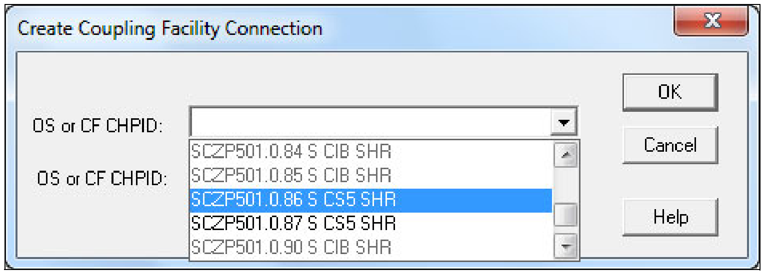

Figure 6-30 HCM - Create Coupling Facility Connection: selecting first CS5 CHPID

19.From the pull-down, select the second CS5 CHPID we defined (SCZP502 - 0.96), and then click **OK**. See Figure 6-31.

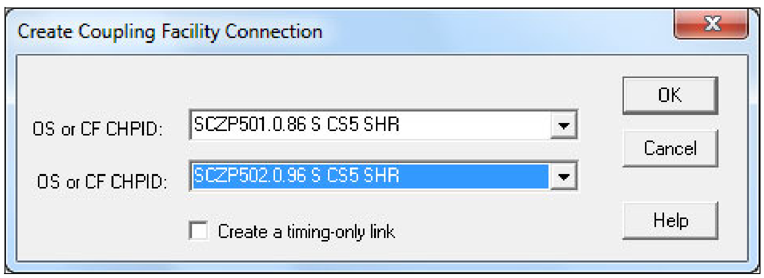

Figure 6-31 HCM - Create Coupling Facility Connection: selecting second CS5 CHPID

20.HCM now suggests control unit and device addresses. Accept or override the control unit and device numbers, then click **OK** (Figure 6-32).

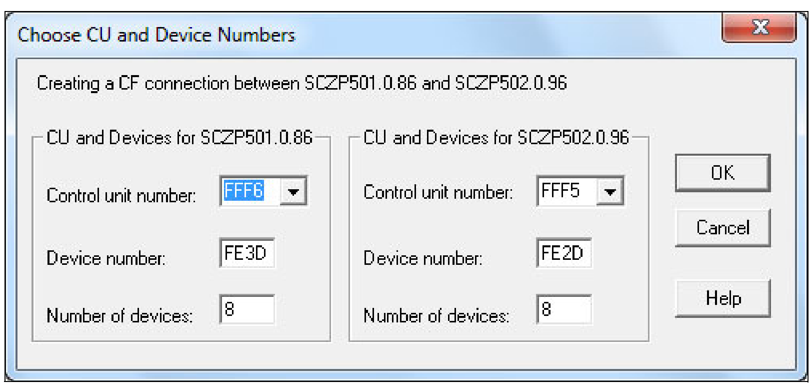

Figure 6-32 HCM - Choose CU and Device Numbers: creating a CF link

21.The CF link is now created. To view it, select the **Locate** → **CF connection** HCM menu option or click one of the CPC CHPID definitions and press F4. See Figure 6-33 on page 280.

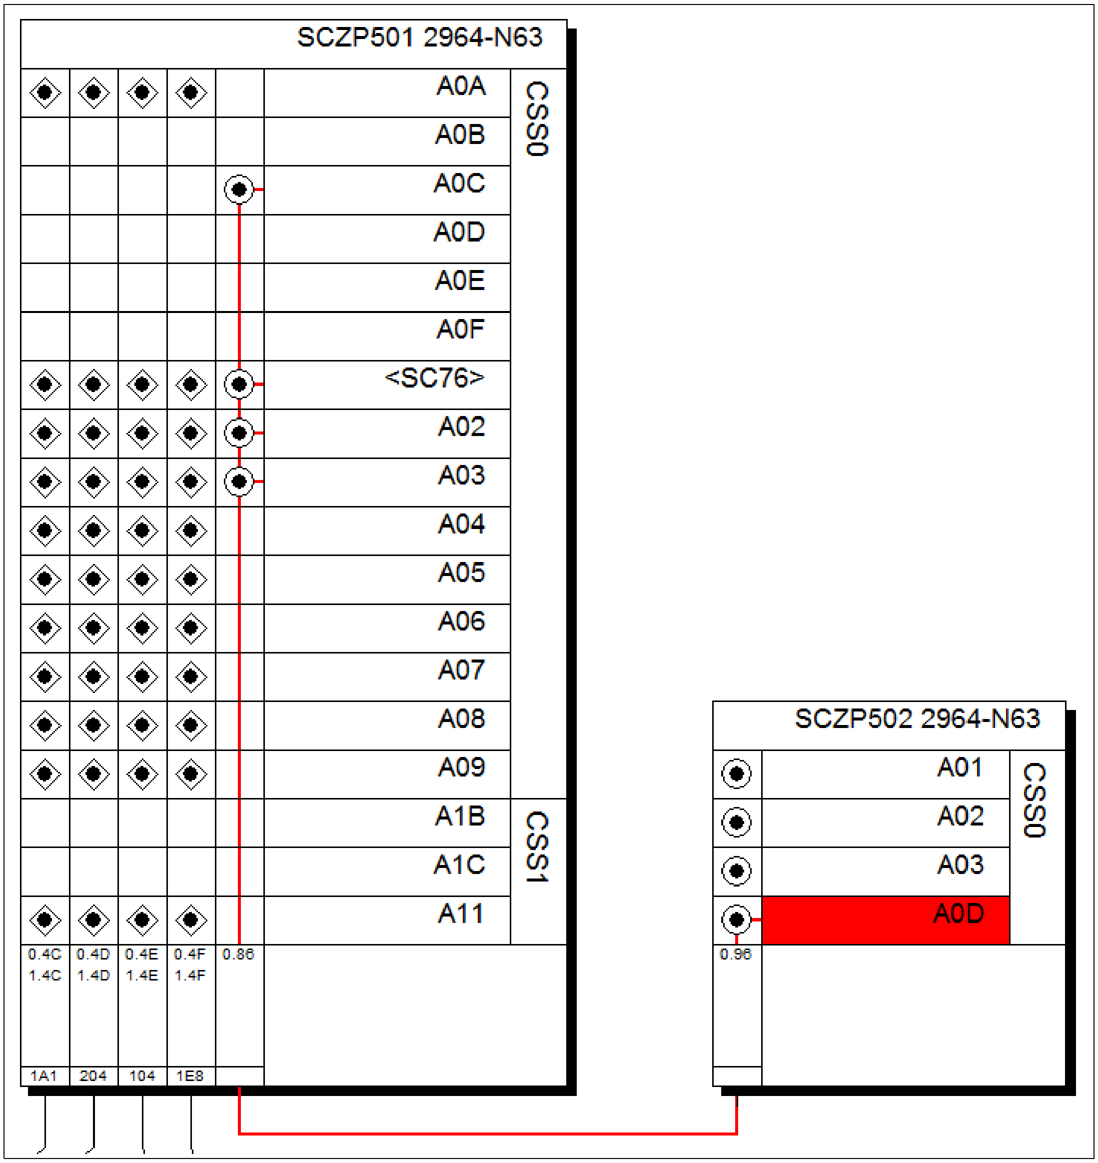

Figure 6-33 HCM - Viewing the CS5 CF link graphically

Defining a RoCE PCIe function using HCD

RDMA over Converged Ethernet (RoCE) uses a new PCIe hardware card for the zEC12 and IBM z13 and is named *10GbE RoCE Express SR* (FC 0411).

To use RoCE hardware functions, Function IDs, Virtual Function IDs, and Physical Network IDs must be defined in HCD to the PCHID that was assigned to the RoCE 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 RoCE PCHIDs.

These steps describe only the definition process. However if you want deeper understanding of how RoCe works, see Chapter 14, “Configure Open System Adaptor (OSA) and RoCE Express” on page 653.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 6-1).

Example 6-1 PCHID Report - RoCE information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                  Comment
19/02/J01       Z22B  20    0411   140                         RG1     NEW
19/12/J01       Z08B  03    0411   208                         RG2     NEW
```

In this example, we define these items:

- PCHID=140 to Function IDs 000 (VF=1) and 001 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1
- PCHID=208 to Function IDs 010 (VF=1) and 011 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1

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 (SCZP502) to which you want to define the RoCE functions and press Enter. See Figure 6-34.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
****************************** Bottom of data *******************************
```

Figure 6-34 HCD - Processor List: adding PCIe functions to a processor

2. Type **add** on the command line in the PCIe Function List panel (Figure 6-35) to add a new PCIe function.

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+          Description
****************************** Bottom of data *******************************
```

Figure 6-35 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 6-36) and press Enter.
   – Function ID: `000`
   – Type: `ROCE`
   – PCHID: `140`
   – Virtual Function ID: `1`
   – Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID  . . . . : SCZP502      test machine definition

Function ID . . . . . . 000
Type  . . . . . . . . . ROCE         +

PCHID . . . . . . . . . 140
Virtual Function ID . . 1_  +

Description . . . . . . RoCE F=000 VF=1 PCHID=140_______
```

Figure 6-36 HCD - PCIe Function List: adding PCIe functions to a processor

4. Enter the following value in the Add/Modify Physical Network IDs panel (Figure 6-37) and press Enter.
   – Physical network ID 1: `ITSOPNET1`

```
-------------------- Add/Modify Physical Network IDs ------------------


If the PCHID is associated to one or more physical networks, specify
each physical network ID corresponding to each applicable physical port

Physical network ID 1  . . ITSOPNET1_______
Physical network ID 2  . . ________________
Physical network ID 3  . . ________________
Physical network ID 4  . . ________________
```

Figure 6-37 HCD - Add/Modify Physical Network IDs: adding network ID

5. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Press Enter. See Figure 6-38.

```
-------------------------- Define Access List ----------------------------
                                                                Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 6-38   HCD - Define Access List: selecting partition for Function access

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 (Figure 6-39) where you can see the Function now defined.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
******************************* Bottom of data ********************************
```

Figure 6-39   HCD - PCIe Function List: Function now created

7. Define the other Function IDs according to the example so far. See Figure 6-40.

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

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
_ 001    140    2    ROCE            RoCE F=000 VF=2 PCHID=140
_ 010    208    1    ROCE            RoCE F=000 VF=1 PCHID=208
_ 011    208    2    ROCE            RoCE F=000 VF=2 PCHID=208
******************************* Bottom of data ********************************
```

Figure 6-40 HCD - PCIe Function List: All functions now created

Defining a RoCE PCIe function using HCM

RDMA over Converged Ethernet (RoCE) and utilizes a new PCIe hardware card for the zEC12 and IBM z13 called *10GbE RoCE Express SR* (FC 0411).

To use RoCE hardware functions, Function IDs, Virtual Function IDs, and Physical Network IDs must be defined in HCD to the PCHID that was assigned to the RoCE 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 RoCE PCHIDs.

These steps describe only the definition process. However, if you want deeper understanding of how RoCe works, see Chapter 14, “Configure Open System Adaptor (OSA) and RoCE Express” on page 653.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 6-2).

Example 6-2 PCHID Report - RoCE information from PCHID report

Source	**Cage**	**Slot**	**F/C**	**PCHID/Ports or AID**		**Comment**
19/02/J01	Z22B	20	0411	140	RG1	NEW
19/12/J01	Z08B	03	0411	208	RG2	NEW

In this example, we define these items:

- PCHID=140 to Function IDs 000 (VF=1) and 001 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1
- PCHID=208 to Function IDs 010 (VF=1) and 011 (VF=2) on CPC = SCZP502 to Physical Network ID 1 = ITSOPNET1

Complete the following steps:

1. From the HCM main page, click **Edit** → **Processor**.
2. Select the first CPC (SCZP502), and then click **OK.**
3. Click **Edit** → **PCIe** (Figure 6-41 on page 285).

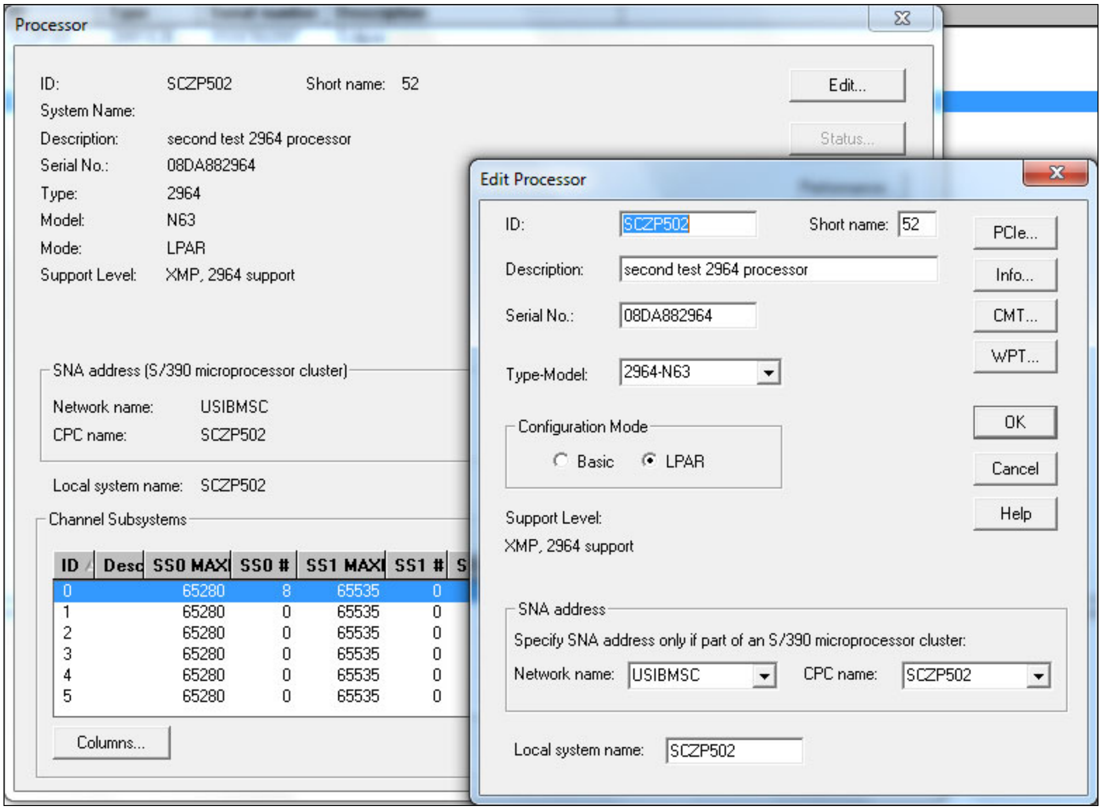

Figure 6-41 HCM - Edit Processor: PCIe functions

4. The PCIe Functions summary window opens (Figure 6-42). Click **Add**.

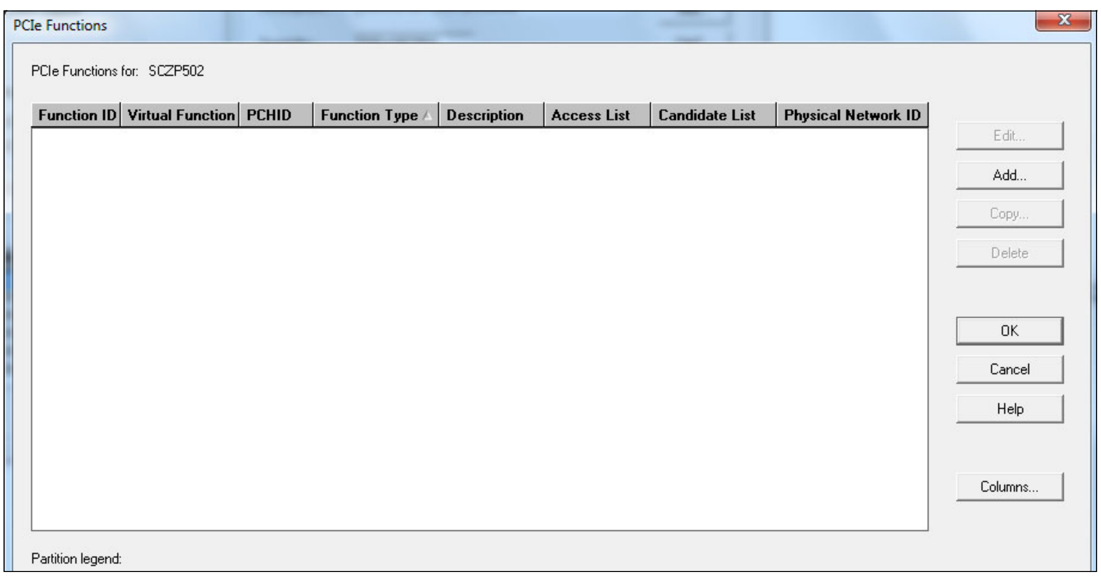

Figure 6-42 HCM - PCIe functions: summary list

5. Enter the following values in the Add PCIe Function window (Figure 6-43 on page 286) and click **OK**.
   - Function ID: `000`
   - Function Type: `ROCE`
   - Virtual Function Number: `1`

– PCHID: `140`
– Description: `optional`
– Physical Network ID 1: `ITSOPNET1`
– LPAR Access List: `0.A01`
– LPAR Candidate List: `0.A02` and `0.A03`

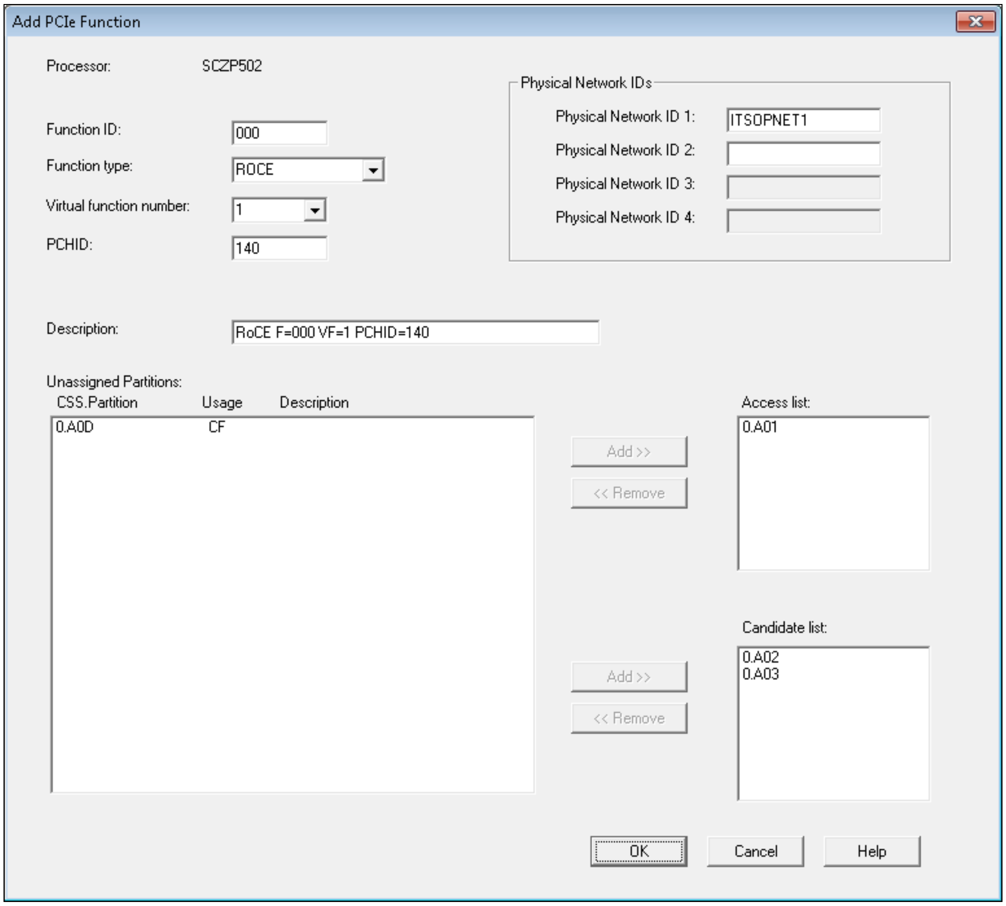

Figure 6-43 HCM - Add PCIe functions: RoCE definition

The RoCE PCIe feature is now listed in the PCIe Function summary list (Figure 6-44).

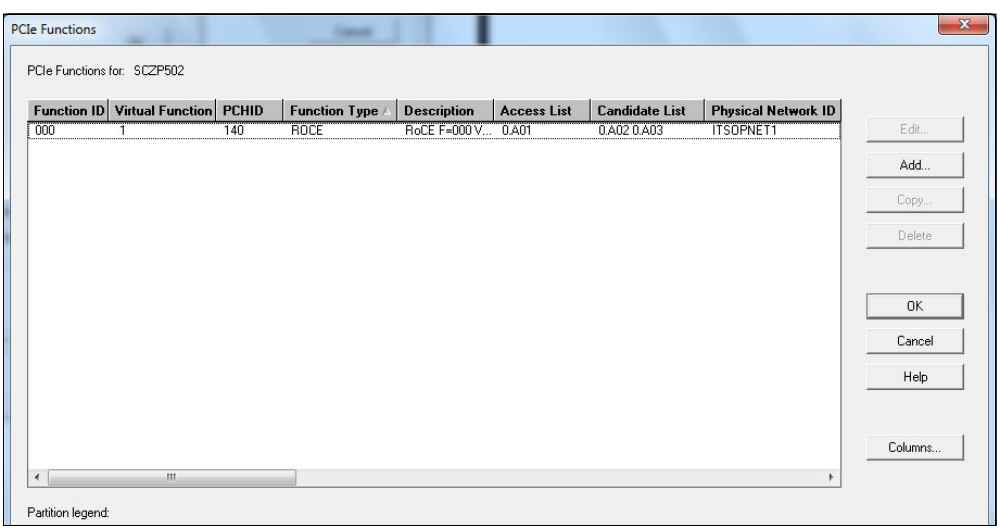

Figure 6-44 HCM - PCIe functions: RoCE PCIe function defined

Defining a zEDC EXPRESS PCIe function using HCD

The zEnterprise Data Compression (zEDC) uses a new PCIe hardware card for the zEC12 and z13 called *zEDC Express FC 0420*.

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that has been assigned to the zEDC hardware cards 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.

Here we show only the definition process. If you want a deeper understanding of how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 6-3).

Example 6-3 PCHID Report - zEDC information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID              Comment
15/14/J01       Z15B  19    0420   1BC                       RG1   NEW
15/05/J01       Z08B  38    0420   27C                       RG2   NEW
```

In this example, we define these items:

- PCHID=1BC to Function IDs 020, 021, and 022 (VF=1, 2, and 3) on CPC = SCZP502
- PCHID=27C to Function IDs 030, 031, and 032 (VF=1, 2, and 3) on CPC = SCZP502

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 (SCZP502) to which you want to define the zEDC functions and press Enter. See Figure 6-45.

```
                              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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
****************************** Bottom of data ********************************
```

Figure 6-45 HCD - Processor List: adding PCIe functions to a processor

2. Type **add** on the command line in the PCIe Function List panel (Figure 6-46 on page 288) to add a new PCIe function.

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 6-46 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 6-47) and press Enter.
   – Function ID: `020`
   – Type: `ZEDC-EXPRESS`
   – PCHID: `1BC`
   – Virtual Function ID: `1`
   – Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID  . . . . : SCZP502      test machine definition

Function ID . . . . . . 020
Type  . . . . . . . . . ZEDC-EXPRESS  +

PCHID . . . . . . . . . 1BC
Virtual Function ID . . 1   +

Description . . . . . . zEDC F=020 VF=1 PCHID=1BC_______
```

Figure 6-47 HCD - PCIe Function List: adding PCIe functions to a processor

4. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Then, press Enter. See Figure 6-48.

```
-------------------------- Define Access List ----------------------------
                                                                Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 6-48 HCD - Define Access List: selecting partition for Function access

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 (Figure 6-49) where you can see that the Function is now defined.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
******************************* Bottom of data ********************************
```

Figure 6-49 HCD - PCIe Function List: Function now created

6. Now define the other Function IDs according to this example. See Figure 6-50.

```
                              PCIe Function List      Row 1 of 6 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
_ 021    1BC    2    ZEDC-EXPRESS    zEDC F=021 VF=2 PCHID=1BC
_ 022    1BC    3    ZEDC-EXPRESS    zEDC F=022 VF=3 PCHID=1BC
_ 030    27C    1    ZEDC-EXPRESS    zEDC F=030 VF=1 PCHID=27C
_ 031    27C    2    ZEDC-EXPRESS    zEDC F=031 VF=2 PCHID=27C
_ 032    27C    3    ZEDC-EXPRESS    zEDC F=032 VF=3 PCHID=27C
****************************** Bottom of data *******************************
```

Figure 6-50 HCD - PCIe Function List: All Functions now created

Defining a zEDC-EXPRESS PCIe function using HCM

zEnterprise Data Compression (zEDC) uses a new PCIe hardware card for the zEC12 and IBM z13 named *zEDC Express FC 0420*.

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that has been assigned to the zEDC hardware cards 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.

Here we show only the definition process. If you want a deeper understanding of how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 6-4).

Example 6-4 PCHID Report - zEDC information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                  Comment
15/14/J01       Z15B  19    0420   1BC                          RG1    NEW
15/05/J01       Z08B  38    0420   27C                          RG2    NEW
```

In this example, we define these items:

- PCHID=1BC to Function IDs 020, 021 and 022 (VF=1, 2 and 3) on CPC = SCZP502
- PCHID=27C to Function IDs 030, 031 and 032 (VF=1, 2 and 3) on CPC = SCZP502

Complete the following steps:

1. From the HCM main page, click **Edit** → **Processor**.
2. Select the first CPC (SCZP502), then click **OK**.
3. Click **Edit** → **PCIe** (Figure 6-51 on page 291).

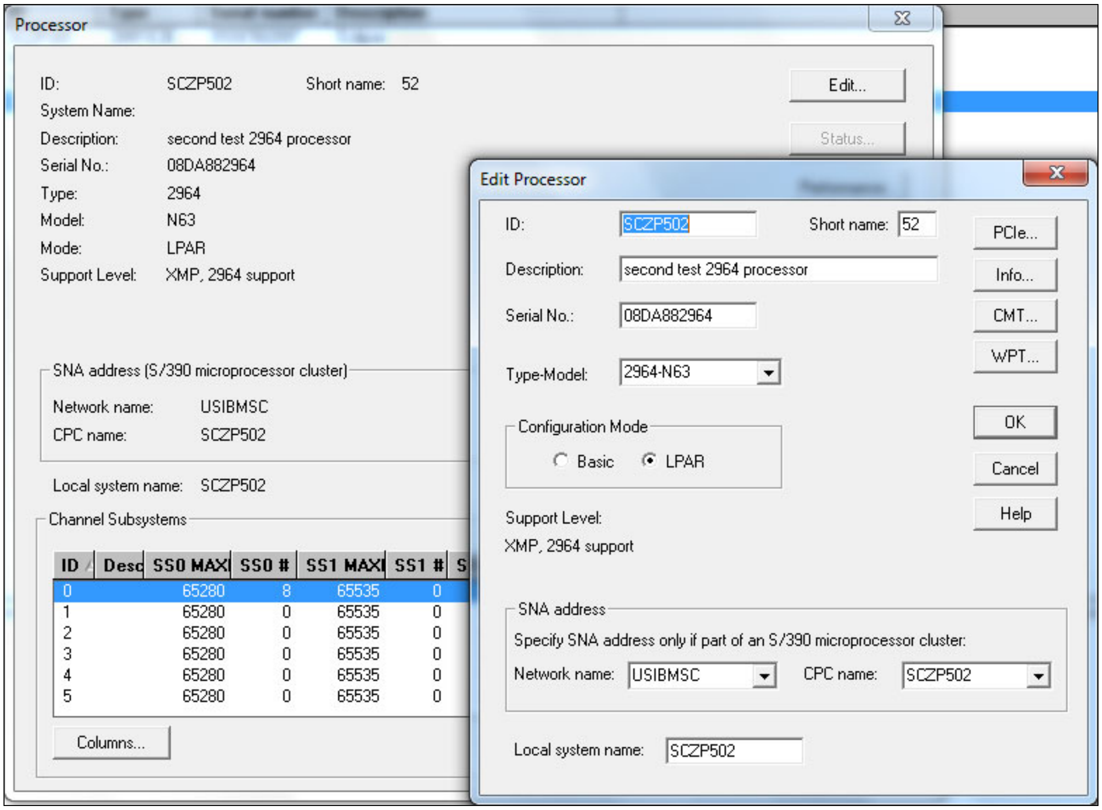

Figure 6-51 HCM - Edit Processor: PCIe functions

4. HCM displays the PCIe Functions summary window (Figure 6-52). Click **Add**.

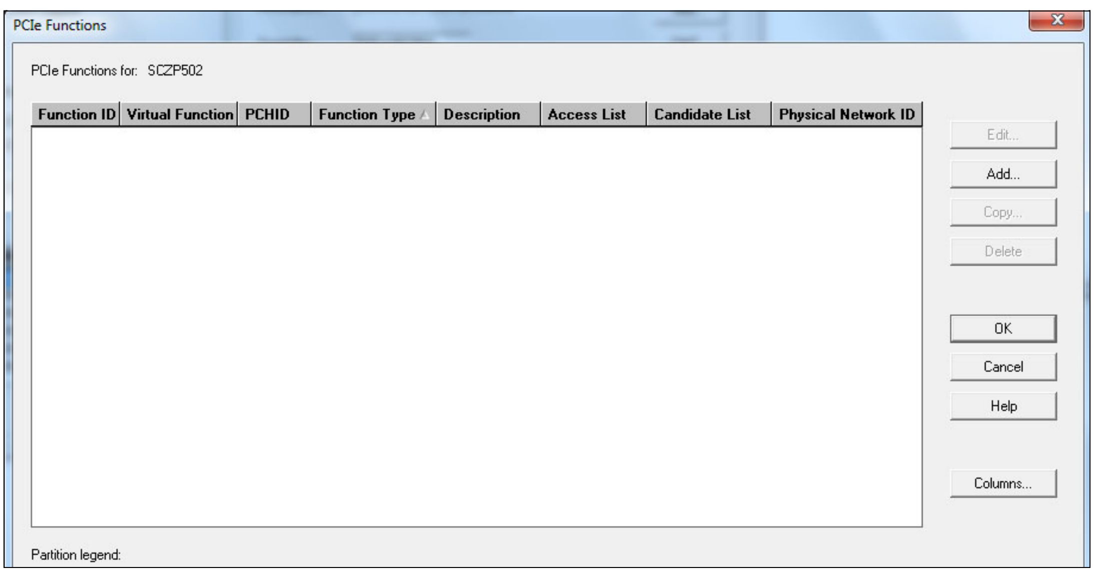

Figure 6-52 HCM - PCIe functions: summary list

5. Enter the following values in the Add PCIe Function window (Figure 6-53 on page 292), and click **OK**.
   - Function ID: `020`
   - Function Type: `ZEDC-EXPRESS`
   - Virtual Function Number: `1`

– PCHID: `1BC`
– Description: `optional`
– LPAR Access List: `0.A01`
– LPAR Candidate List: `0.A02` and `0.A03`

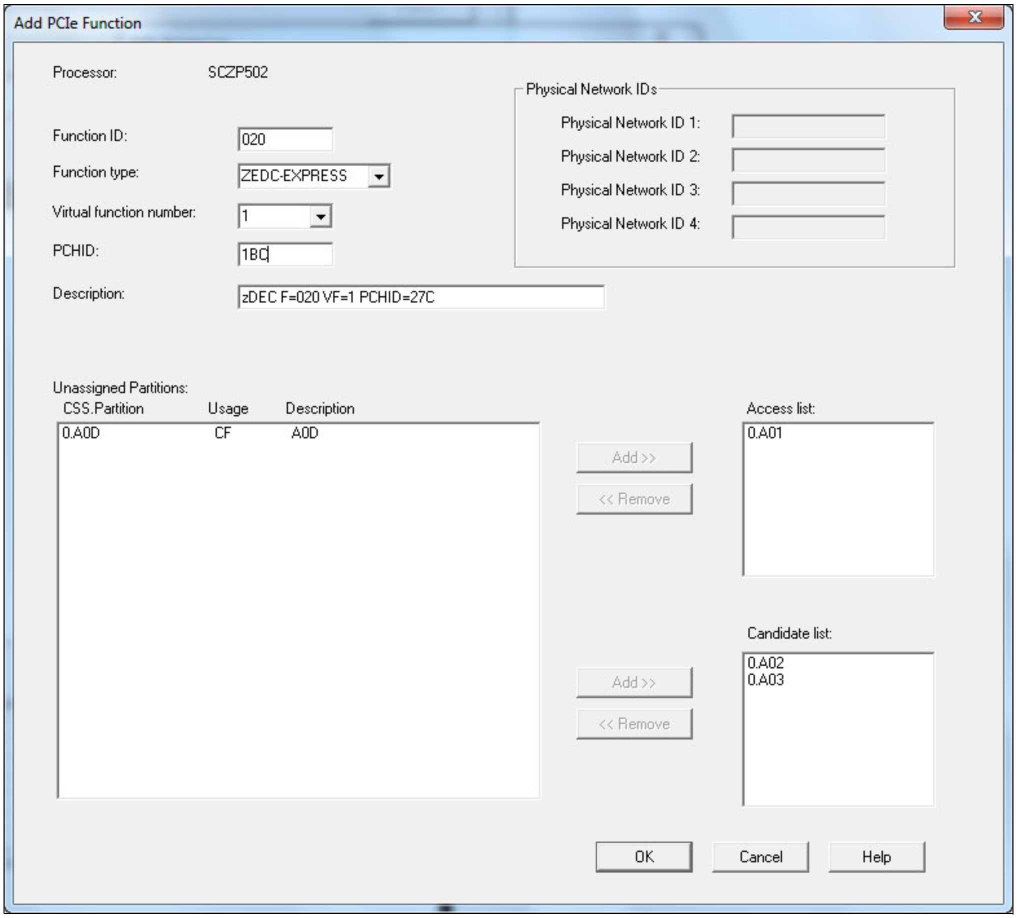

Figure 6-53 HCM - Add PCIe functions: zEDC definition

The zEDC PCIe feature is now listed in the PCIe Function summary list (Figure 6-54).

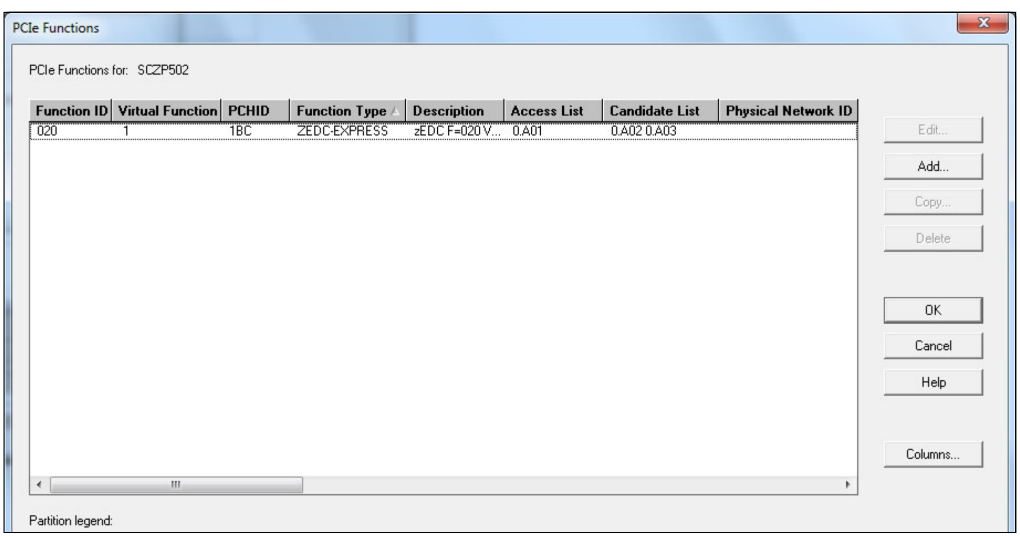

Figure 6-54 HCM - PCIe functions: zEDC PCIe function defined

6.2.8 Overdefining channel paths on an XMP processor

Sometimes you must define a channel path that is not physically installed on the processor. This definition is useful if you are planning to add more channel cards to the processor in the future and want to have the definitions in the IODF before the hardware is installed.

HCD allows you to overdefine CHPIDs by using an asterisk (*) for the PCHID value. An overdefined CHPID must adhere to all validation rules, but it is not taken into account by an IOCDS download. Also, it is not included in the IOCP statements, in a CONFIGxx member, or during dynamic activation.

If a control unit contains only CHPIDs with a PCHID value of an asterisk (*), the whole control unit (including any attached devices) is omitted from the configuration to be activated.

When you install the channel path later, you must edit the CHPID and replace the asterisk (*) with its valid PCHID.

> **Remember:** This is not the case for CIB type CHPIDs, where these CHPIDs have connections to other CIB type CHPIDs. Therefore, HCD allows you to define CIB type CHPIDs as overdefined only if they are unconnected.
>
> Overdefining is also supported for CS5 type CHPID definitions.

The 2964 production IODF can then be activated dynamically and the PCHID, CHPID, and control unit definitions become available to the operating system.

Figure 6-55 shows what the CHPID/PCHID definitions look like before they are set as overdefined. Press PF20 (right) in the Channel Path List.

```
------------------------------------------------------------------------------
                              Channel Path List    Row 83 of 105 More: <     >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP401    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 81    18/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 8B    1C/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 8C    1C/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 90    08/1  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ A2    210   CFP   SHR   No    ________     _ _ _ _ a _ _ _ _ _ _ _ a _ _
_ A3    198   CFP   SHR   No    ________     _ _ _ _ a _ _ _ _ _ _ _ a _ _
_ C0    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ _ a
_ C1    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ a _
_ C2    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ _ a
_ C3    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ a _
```

Figure 6-55 HCD - Channel Path List: reserving CHPIDs

Figure 6-56 shows what the CHPID/PCHID definitions look like after they are set as overdefined.

```
-------------------------------------------------------------------------------
                              Channel Path List    Row 83 of 105 More: <     >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP401    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 81    18/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 8B    1C/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 8C    1C/2  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ 90    08/1  CIB   SHR   No    ________     _ _ _ _ _ _ _ _ _ _ a _ _ _ _
_ A2    *     CFP   SHR   No    ________     _ _ _ _ a _ _ _ _ _ _ _ a _ _
_ A3    *     CFP   SHR   No    ________     _ _ _ _ a _ _ _ _ _ _ _ a _ _
_ C0    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ _ a
_ C1    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ a _
_ C2    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ _ a
_ C3    ____  ICP   SHR   No    ________     a _ a _ _ _ _ _ _ _ _ _ _ a _
```

Figure 6-56 HCD - Channel Path List: overdefined CHPIDs

6.3 OSA: Saving and restoring configuration data

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

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

See 14.1, “Configure OSE channel with Open System Facility (OSA/SF) on the HMC” on page 654 for details of how to save and restore any Open Systems Adapter (OSA) configuration data such the OSA Address Table (OAT).

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

See 14.5, “Configure OSA-Express Integrated Console Controller (OSA-ICC)” on page 686 if you are unfamiliar with the exporting and importing process for OSA-ICC Server and Session configuration data.

6.3.3 Using OSA Advanced Facilities to set OSA parameters

See 14.6, “Setting OSA parameters by using OSA Advanced Facilities” on page 702 for the process on how to change the OSA port speed or MAC addresses.

6.4 HCD: Validating the 2964 work IODF

This section describes the steps to validate the 2964 work IODF.

To validate the work IODF by using 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` (Figure 6-57).

```
--------------- View I/O Definition File Information -----------


IODF name  . . . . . . . : 'SYS6.IODFB6.WORK'
IODF type  . . . . . . . : Work - Validated
IODF version . . . . . . : 5

Creation date  . . . . . : 2014-11-07
Last update  . . . . . . : 2014-11-10  10:18

Volume serial number . . : IODFPK
Allocated space  . . . . : 4000    (Number of 4K blocks)
Used space . . . . . . . : 2188    (Number of 4K blocks)
   thereof utilized (%)  84
Activity logging . . . . : No
Multi-user access  . . . : No
Backup IODF name . . . . :

Description  . . . . . . :
```

Figure 6-57 HCD - View I/O Definition File Information: validated work IODF

6.5 Creating the IOCP for the CHPID Mapping Tool

To create the IOCP file for the CHPID Mapping Tool, complete the following steps:

> **Tip:** You might prefer to use HCM to create the IOCP statements file and transfer the file to your workstation. You can then start the CHPID Mapping Tool, create an updated IOCP statements file, and transfer the file back to the host.

1. Select HCD option **3. Build IOCP input data set** and press Enter (Figure 6-58).

```
----- 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
```

Figure 6-58 HCD - Activate or Process Configuration Data: building IOCP for SCZP501

2. HCD displays the list of available processors (Figure 6-59). Select the SCZP501 processor by typing a forward slash (/) next to it and pressing Enter.

```
------------------------- Available Processors --------------------------
                                                                Row 1 of 2
Command ===> __________________________________________________________

Select one.

  Processor ID  Type     Model    Mode  Description
  SCZP201       2097     E26      LPAR  Eclipse
/ SCZP501       2964     N63      LPAR  Sphinx
**************************** Bottom of data *****************************
```

Figure 6-59 HCD - 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 (Figure 6-60). Complete the following fields:
   - Title1: `IODFB6`
   - IOCP input data set: `'SYS6.IODFB6.IOCPIN.SCZP501'`
   - 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 . . . . . . . . . : 'SYS6.IODFB6.WORK'
Processor ID  . . . . . . . : SCZP501
Title1 . IODFB6_________________________________________________________
Title2 : SYS6.IODFB6.WORK - 2014-11-10 10:18

IOCP input data set
'SYS6.IODFB6.IOCPIN.SCZP501'__________________________
Input to Stand-alone IOCP?  Yes  (Yes or No)

Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 6-60   HCD - 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 (Figure 6-61). This data set is used as inputto the CHPID Mapping Tool.

```
ID    MSG1='IODFB6',                                              *
      MSG2='SYS6.IODFB6.WORK - 2014-11-10 10:18',                 *
      SYSTEM=(2964,1),LSYSTEM=SCZP501,                            *
      TOK=('SCZP501',00800001DA87296410184795011431 4F00000000,*
      00000000,'14-11-10','10:18:47','........','........')
RESOURCE PARTITION=((CSS(0),(A0A,A),(A0B,B),(A0C,C),(A0D,D),(A*
      0E,E),(A0F,F),(A01,1),(A02,2),(A03,3),(A04,4),(A05,5),(A*
      06,6),(A07,7),(A08,8),(A09,9)),(CSS(1),(A1A,A),(A1B,B),(*
      A1C,C),(A1D,D),(A1E,E),(A1F,F),(A11,1),(A12,2),(A13,3),(*
      A14,4),(A15,5),(A16,6),(A17,7),(A18,8),(A19,9)),(CSS(2),*
      (A2A,A),(A2B,B),(A2C,C),(A2D,D),(A2E,E),(A2F,F),(A21,1),*
      (A22,2),(A23,3),(A24,4),(A25,5),(A26,6),(A27,7),(A28,8),*
      (A29,9)),(CSS(3),(A3A,A),(A3B,B),(A3C,C),(A3D,D),(A3E,E)*
      ,(A3F,F),(A31,1),(A32,2),(A33,3),(A34,4),(A35,5),(A36,6)*
      ,(A37,7),(A38,8),(A39,9)),(CSS(4),(*,1),(*,2),(*,3),(*,4*
      ),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),*
      (*,E),(*,F)),(CSS(5),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6)*
      ,(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)))
FUNCTION FID=000,VF=1,PCHID=140,PNETID=ITSOPNET1,                 *
      PART=((A15),(=)),TYPE=ROCE
FUNCTION FID=010,VF=1,PCHID=208,PNETID=ITSOPNET1,                 *
      PART=((A03),(=)),TYPE=ROCE
FUNCTION FID=020,VF=1,PCHID=1BC,PART=((A15),(=)),TYPE=ZEDC
FUNCTION FID=021,VF=2,PCHID=1BC,PART=((A03),(=)),TYPE=ZEDC
```

Figure 6-61 HCD - IOCP input data set: contents (truncated)

Part of the TOK statement is now replaced with dots (Example 6-5).

Example 6-5 HCD - IOCP file: TOK statement

```
TOK=('SCZP501',00800001DA872964101847950114314F00000000,*
00000000,'14-11-10','10:18:47','........','........')
```

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 only from production IODFs.

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 file from TSO to your 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 `SCZP501in.iocp`.

6.6 CMT: Assigning PCHIDs to CHPIDs

The following steps use the output from the previous set of HCD steps (IOCP) and the 2964 order process (CFReport). Use the CHPID Mapping Tool (CMT) to assign PCHIDs to each of the CHPIDs for the 2964.

For this process, the CMT must be downloaded. For more information about downloading and installing the CMT, see 3.1.5, “CHPID Mapping Tool (CMT)” on page 44. If CMT is already installed, verify that the latest updates are installed. For more information, see the *CHPID Mapping Tool User’s Guide*, GC28-6947-00.

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 PCHID conflict.
4. Process the hardware resolution.
5. Resolve manually the CIB 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 and transfer it back to the host z/OS image.

Requirement: When you upgrade from a 2817 or a 2827 to a 2964, you must use the CHPID Mapping Tool level that supports the 2964.

6.6.1 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. Right-click the Projects pane and select **New** → **Standard CMT Project** (Figure 6-62).

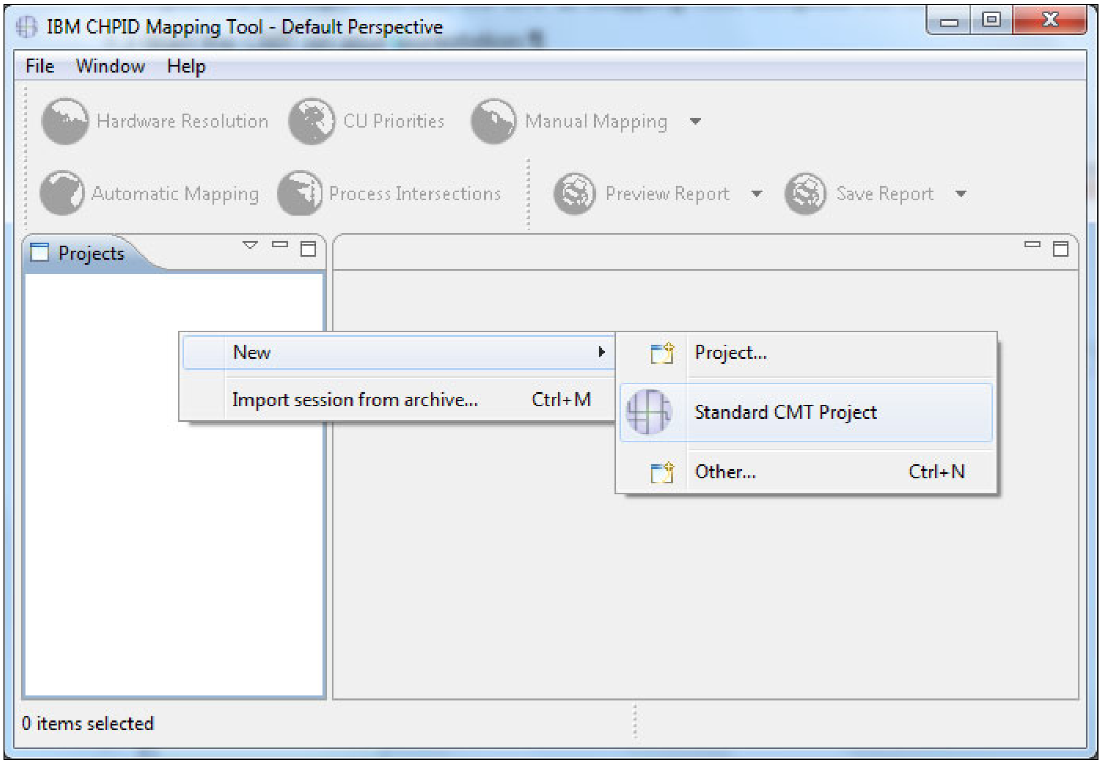

Figure 6-62 CMT - Creating a new CHPID Mapping Tool Project

3. The New CHPID Mapping Tool Project window opens (Figure 6-63). Use this window to specify a project name and then click **Next**.

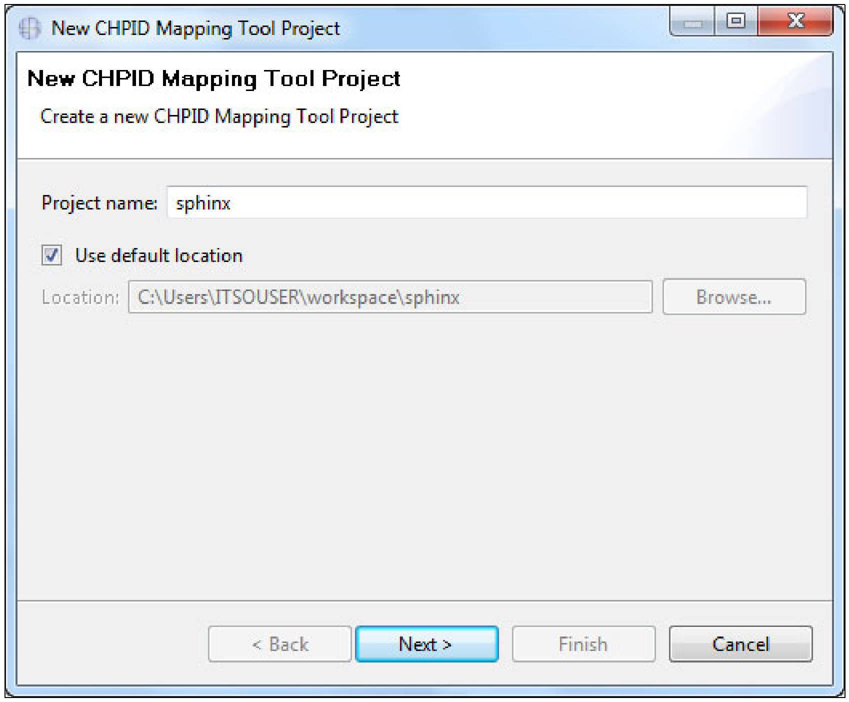

Figure 6-63 CMT - New CHPID Mapping Tool Project name

4. Import the `CFReport` file into the CHPID Mapping Tool by specifying the name in the CFReport file field, and then click **Finish** (Figure 6-64).

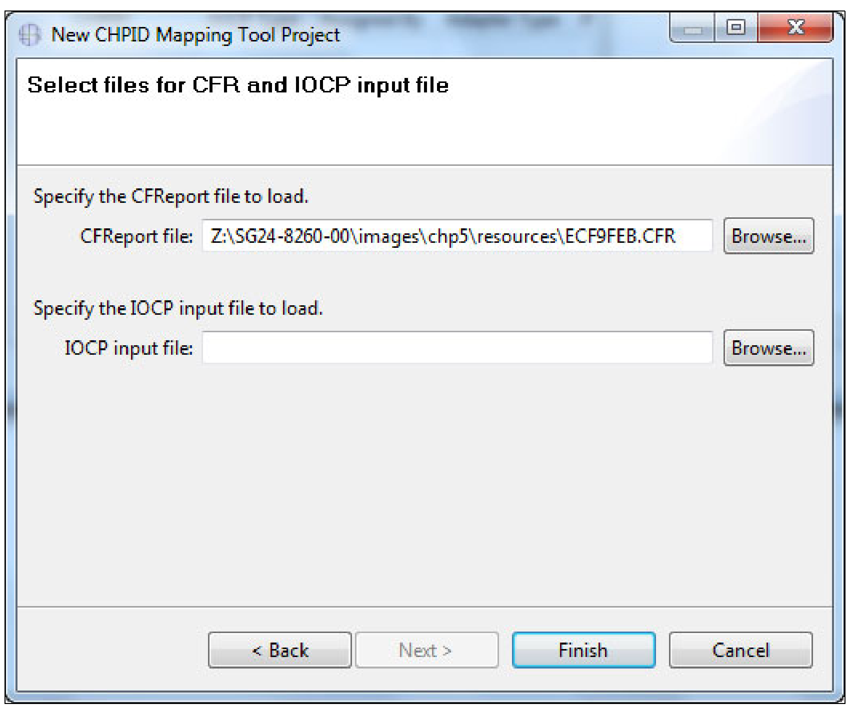

Figure 6-64 CMT - Specifying the CFReport file

If you click **Finish** but did not select an IOCP file, you receive the message shown in Figure 6-65.

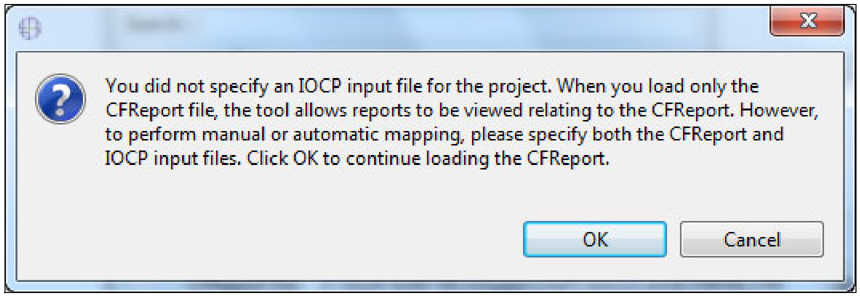

Figure 6-65 CMT - Warning message for not specifying an IOCP file

A window shows the progress of reading the CFReport file (Figure 6-66).

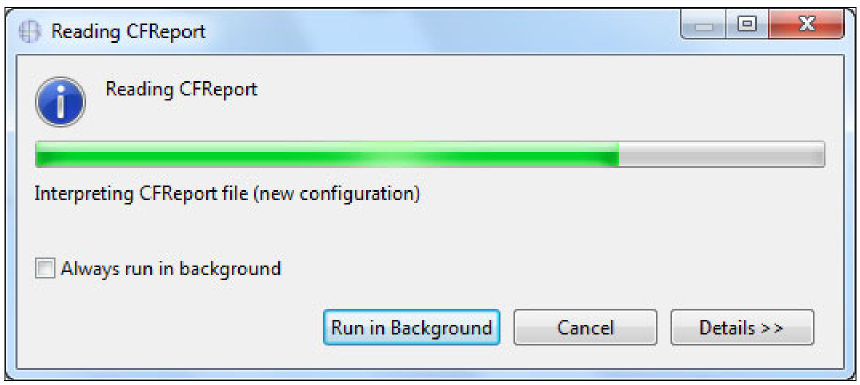

Figure 6-66 CMT - Reading the CFReport file

The information from the CFReport file is displayed in the Hardware pane (Figure 6-67).

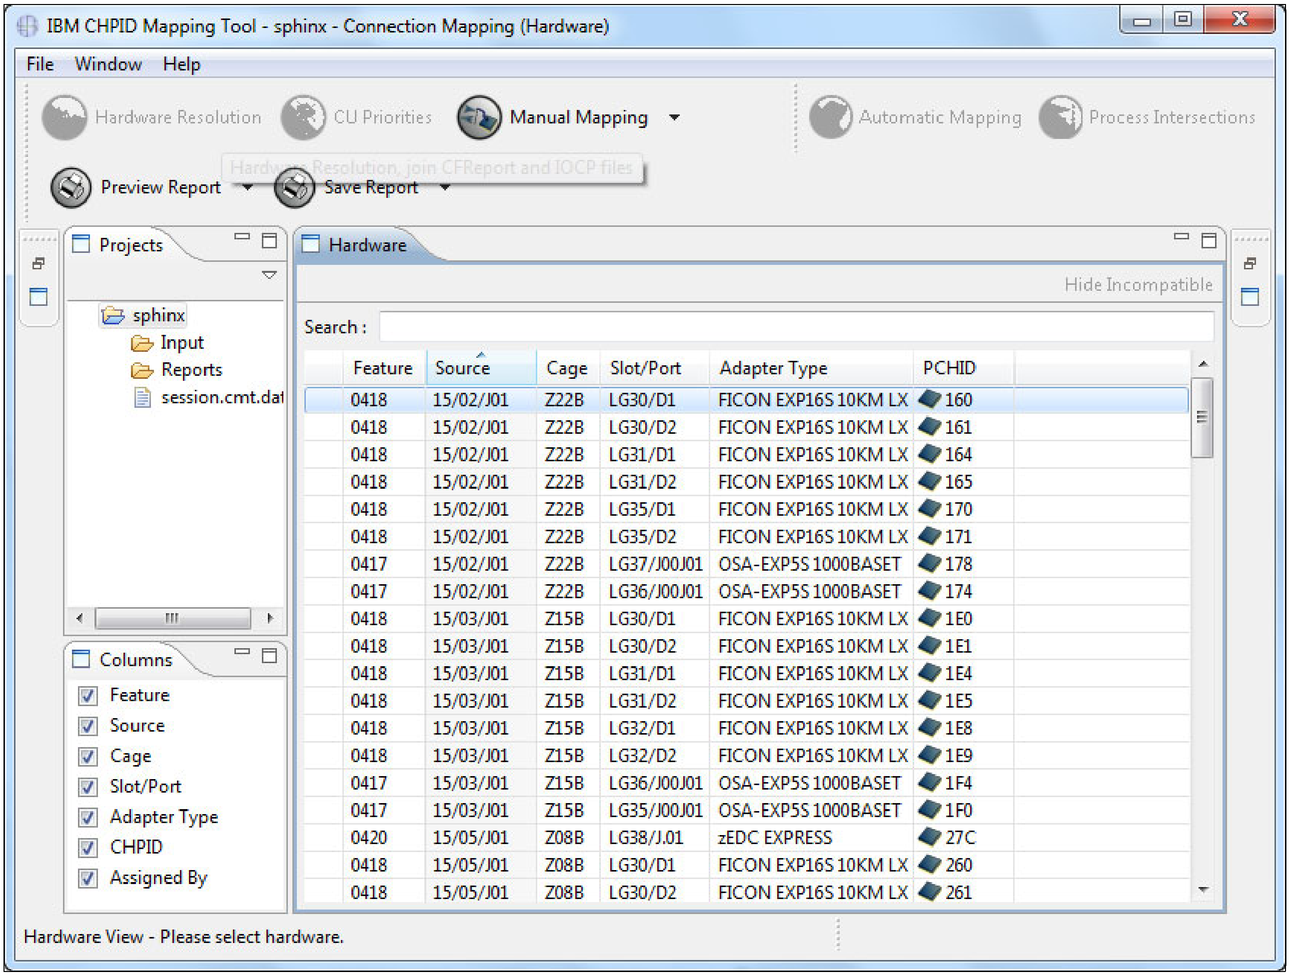

Figure 6-67 CMT - Imported CFReport file

6.6.2 Importing the 2964 IOCP file into the CHPID Mapping Tool

To import the 2964 IOCP file into the CHPID Mapping Tool, complete the following steps:

1. Right-click in the Projects pane and select **Import IOCP input file** (Figure 6-68).

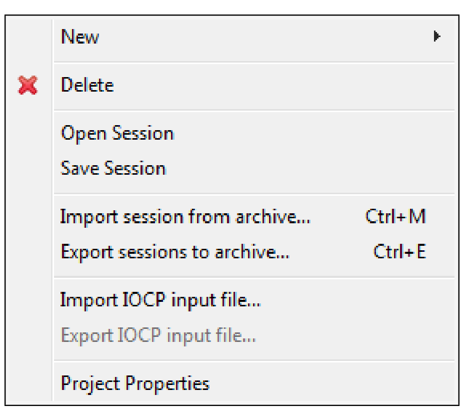

Figure 6-68 CMT - Importing the IOCP file

2. Select the IOCP file on your workstation to import into the CHPID Mapping Tool and click **Finish** (Figure 6-69).

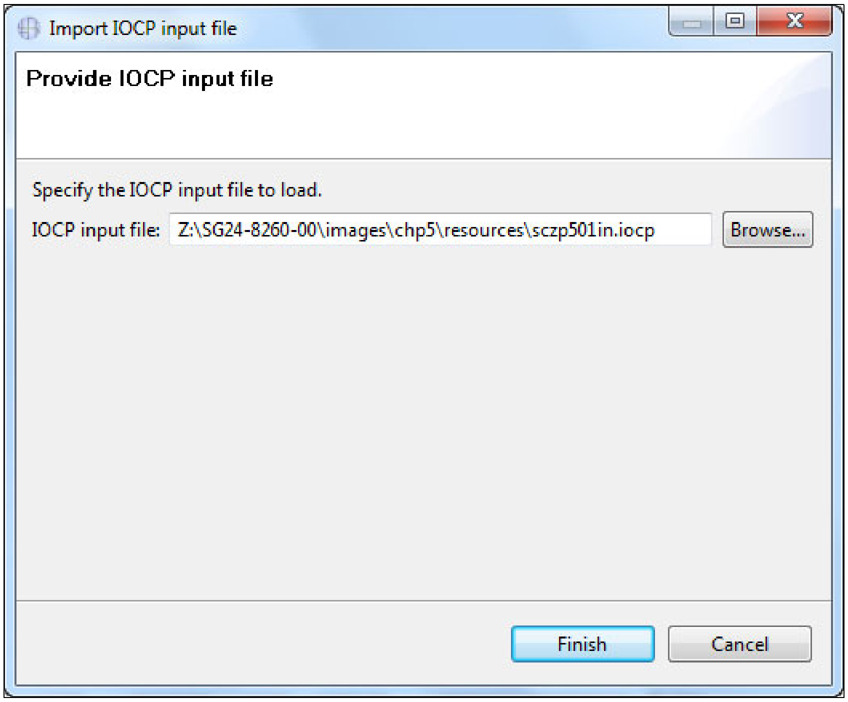

Figure 6-69 CMT - Specifying the IOCP file for import

3. In the Projects pane, under the Input tab, expand the IOCP tab, right-click on the IOCP file, and select **Read Selected IOCP** (Figure 6-70).

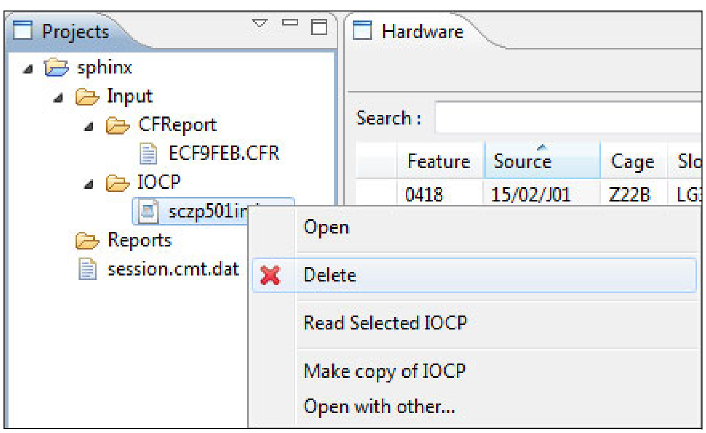

Figure 6-70 CMT - Reading the selected IOCP

A window shows the progress information (Figure 6-71).

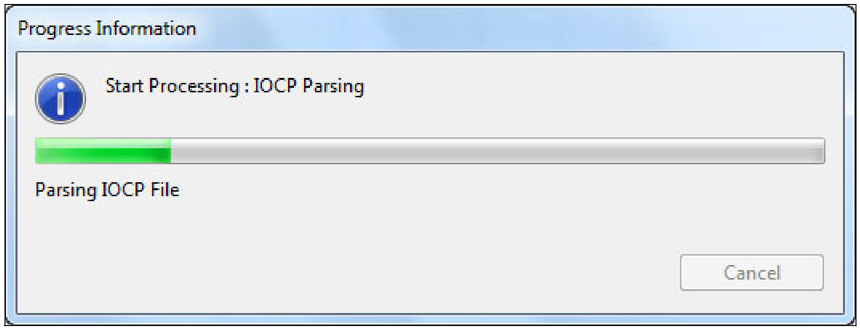

Figure 6-71 CMT - Processing the IOCP file

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 (Figure 6-72) includes three tabbed panes: Projects, Hardware Resolution, and Adapter Type Summary. Hardware Resolution is the middle pane and the Adapter Type Summary is on the right.

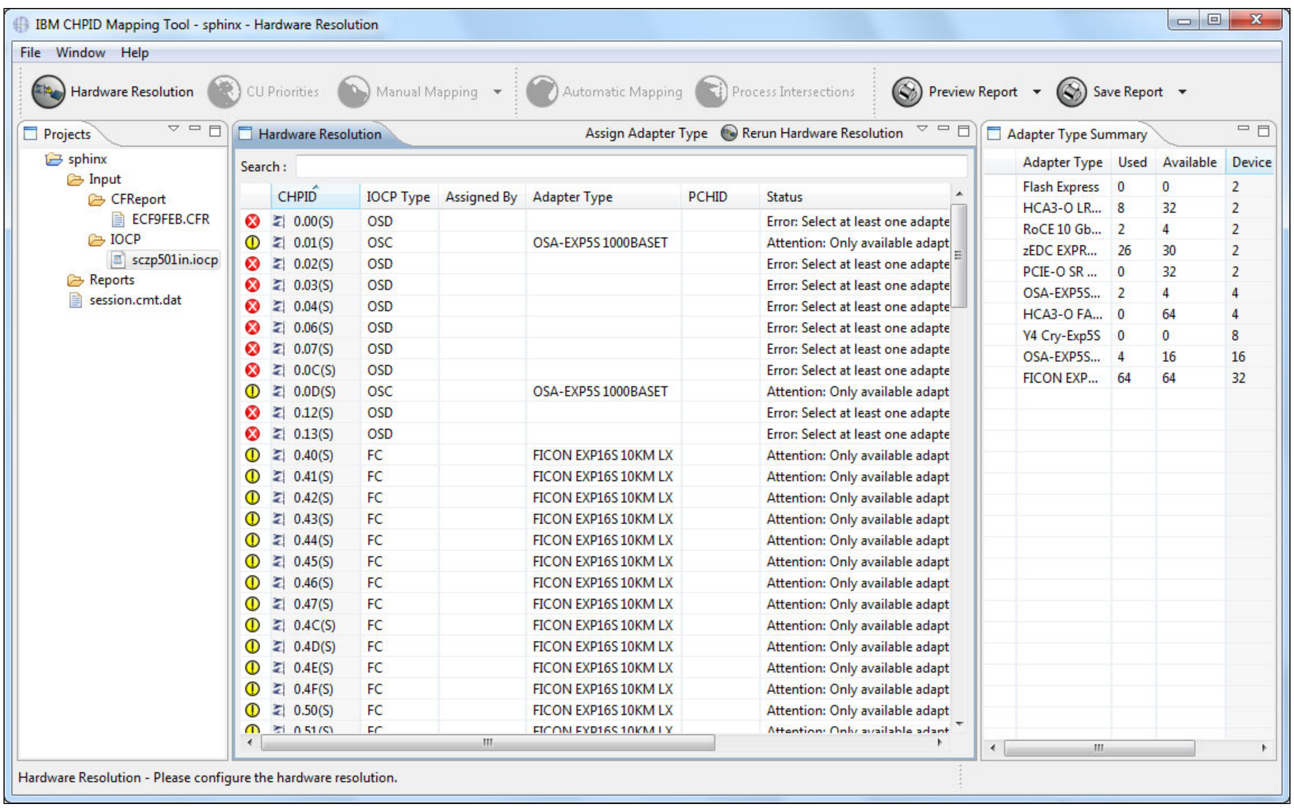

Figure 6-72 CMT - Hardware Resolution after Imported CFReport file

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 shows the following output:

- Hardware Resolution: This window lists all CHPIDs that were found; the Status column shows the CHPID information to be investigated. In the example, investigate the status; Table 6-3 lists status messages and possible resolutions.

Table 6-3 Status messages and possible resolutions

Status	Explanation	Resolution (if required)
No hardware found	AID values or PCHID values are present that are not found in the hardware. This situation could occur when you are replacing hardware for an MES and the IOCP file contains a PCHID value for the old hardware. (The IOCP file contains a PCHID value for the hardware being removed.)	If you have any CHPIDs of IOCP type CIB, 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. Do the following steps for CIB CHPIDs: 1. Remove the AID values. 2. Do one of the following tasks: - Inside the CHPID Mapping Tool, perform manual mapping to associate these CHPIDs with AIDs. - Assign the AID values outside the tool, for example, using Hardware Configuration Definition (HCD). 3. Replace the IOCP file.
Select at least one adapter type.	A adapter type is not assigned to the current row.	Assign a adapter type to IOCP type.
Adapter_type is not compatible with *IOCP_type*.	Adapter type assigned for the CHPID is not compatible with the IOCP type specified by the IOCP file.	See Performing hardware resolution for a type mismatch
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 need to change the IOCP file or obtain additional hardware
PCHID_1 moved to new channel ID: PCHID_2 **Example**: 520 moved to 1E2	You are replacing hardware for an MES, and the IOCP file contains a PCHID value for the old hardware, which is being removed. This PCHID value has moved from an old machine to the PCHID value for the new hardware. PCHID_1 is the first PCHID value (for example, 520) and PCHID_2 is the second PCHID 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 if you prefer a different assignment.

- Manual mapping CIB CHPIDs: Availability Mapping cannot be used until all CIB CHPIDs are resolved. You can use manual mapping to resolve any CIB 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 PCHID in the Manual window. If this option is not selected (it has no check mark), then availability PCHIDs 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 might require recabling after availability assignments.

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

    Regardless of whether the channels are moving or not, the CHPID Mapping Tool either assigns PCHIDs 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 PCHID 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 only on CHPIDs that do not have PCHIDs assigned.

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

  > **Attention:** If you run **Reset CHPIDs assigned by IOCP**, it will reset any previously mapped CHPID assignments and 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 prior CHPID-to-PCHID relationship does not cause unacceptable availability.

6.6.3 Resolving CHPIDs with PCHID conflict

The CMT displays the CHPIDs with PCHID conflicts (Figure 6-73).

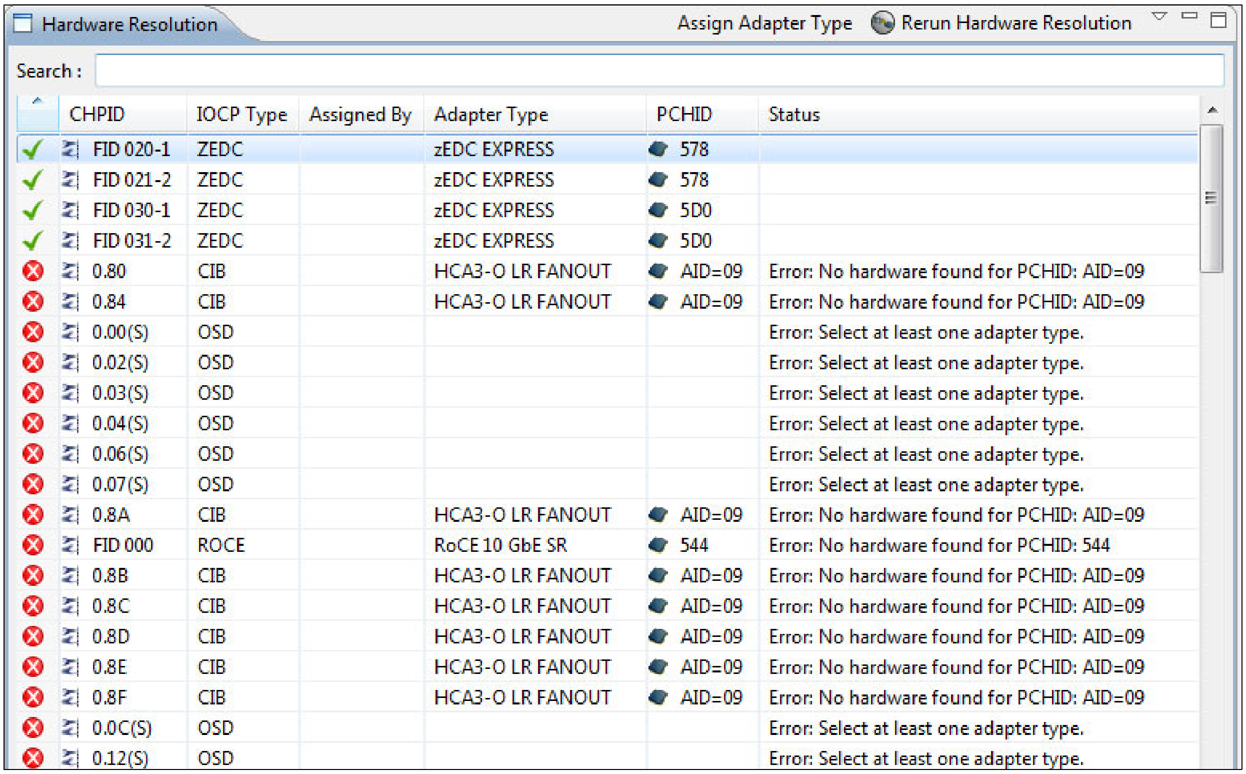

	CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
	FID 020-1	ZEDC		zEDC EXPRESS	578	
	FID 021-2	ZEDC		zEDC EXPRESS	578	
	FID 030-1	ZEDC		zEDC EXPRESS	5D0	
	FID 031-2	ZEDC		zEDC EXPRESS	5D0	
	0.80	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.84	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.00(S)	OSD				Error: Select at least one adapter type.
	0.02(S)	OSD				Error: Select at least one adapter type.
	0.03(S)	OSD				Error: Select at least one adapter type.
	0.04(S)	OSD				Error: Select at least one adapter type.
	0.06(S)	OSD				Error: Select at least one adapter type.
	0.07(S)	OSD				Error: Select at least one adapter type.
	0.8A	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	FID 000	ROCE		RoCE 10 GbE SR	544	Error: No hardware found for PCHID: 544
	0.8B	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.8C	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.8D	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.8E	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.8F	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
	0.0C(S)	OSD				Error: Select at least one adapter type.
	0.12(S)	OSD				Error: Select at least one adapter type.

Figure 6-73 CMT - CHPIDs with PCHID conflicts

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

- An X in a red circle: This indicates an error.
- A green check mark: This indicates that the tool successfully resolved the specified Channel Type.

The example has these reasons to resolve hardware resolution issues:

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

6.6.4 Hardware resolution

In the example, the CHPID Mapping Tool displays an X in the first column of the Hardware Resolution pane (Figure 6-74) and is related to these error types: `No hardware found` and `FICON EXP8S 10KM LX is not compatible with OSD`.

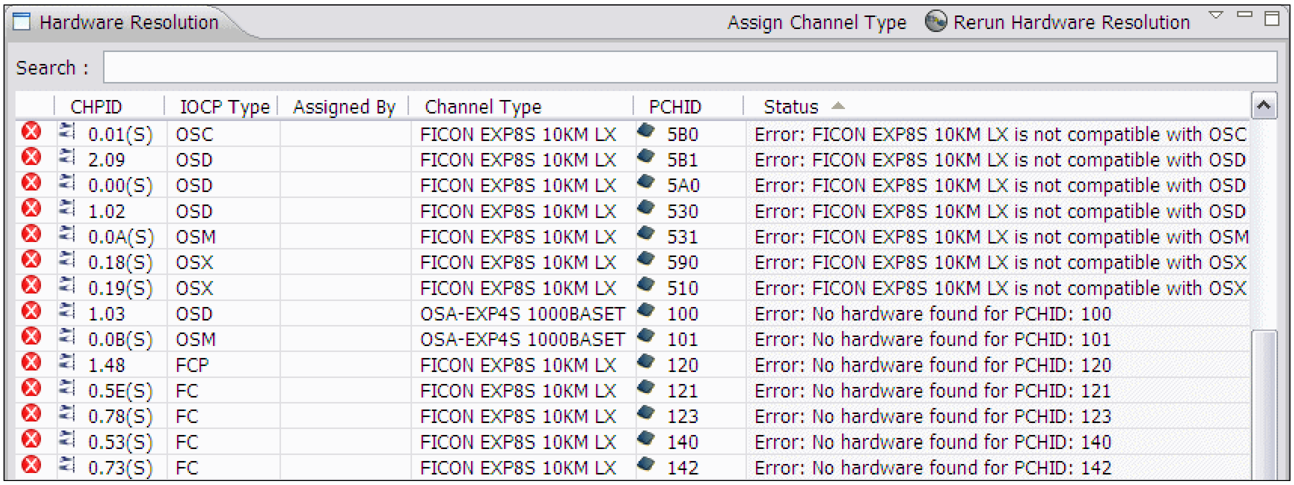

CHPID	IOCP Type	Assigned By	Channel Type	PCHID	Status
0.01(S)	OSC		FICON EXP8S 10KM LX	5B0	Error: FICON EXP8S 10KM LX is not compatible with OSC
2.09	OSD		FICON EXP8S 10KM LX	5B1	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.00(S)	OSD		FICON EXP8S 10KM LX	5A0	Error: FICON EXP8S 10KM LX is not compatible with OSD
1.02	OSD		FICON EXP8S 10KM LX	530	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.0A(S)	OSM		FICON EXP8S 10KM LX	531	Error: FICON EXP8S 10KM LX is not compatible with OSM
0.18(S)	OSX		FICON EXP8S 10KM LX	590	Error: FICON EXP8S 10KM LX is not compatible with OSX
0.19(S)	OSX		FICON EXP8S 10KM LX	510	Error: FICON EXP8S 10KM LX is not compatible with OSX
1.03	OSD		OSA-EXP4S 1000BASET	100	Error: No hardware found for PCHID: 100
0.0B(S)	OSM		OSA-EXP4S 1000BASET	101	Error: No hardware found for PCHID: 101
1.48	FCP		FICON EXP8S 10KM LX	120	Error: No hardware found for PCHID: 120
0.5E(S)	FC		FICON EXP8S 10KM LX	121	Error: No hardware found for PCHID: 121
0.78(S)	FC		FICON EXP8S 10KM LX	123	Error: No hardware found for PCHID: 123
0.53(S)	FC		FICON EXP8S 10KM LX	140	Error: No hardware found for PCHID: 140
0.73(S)	FC		FICON EXP8S 10KM LX	142	Error: No hardware found for PCHID: 142

Figure 6-74 CMT - Hardware resolution status errors

> **More information:** For more information about these error messages, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

The options that must be reset are as follows:

- Reset Incompatible (Hardware - I/O) Entries: (shown in example)
- Reset "Error: No hardware found" Entries: (shown in example)
- Reset "Select at least one adapter type": (shown in example)
- Reset "Required hardware for type IOCP_type not available": (*not* shown in example)
- Reset "PCHID_1 moved to new channel ID: PCHID_2": (*not* shown in example)

Reset Incompatible (Hardware - I/O) Entries

The Channel type that is assigned for the CHPID is not compatible with the IOCP type 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 PCHID. In the example, the IOCP type is OSD but the PCHID 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 (Figure 6-75).

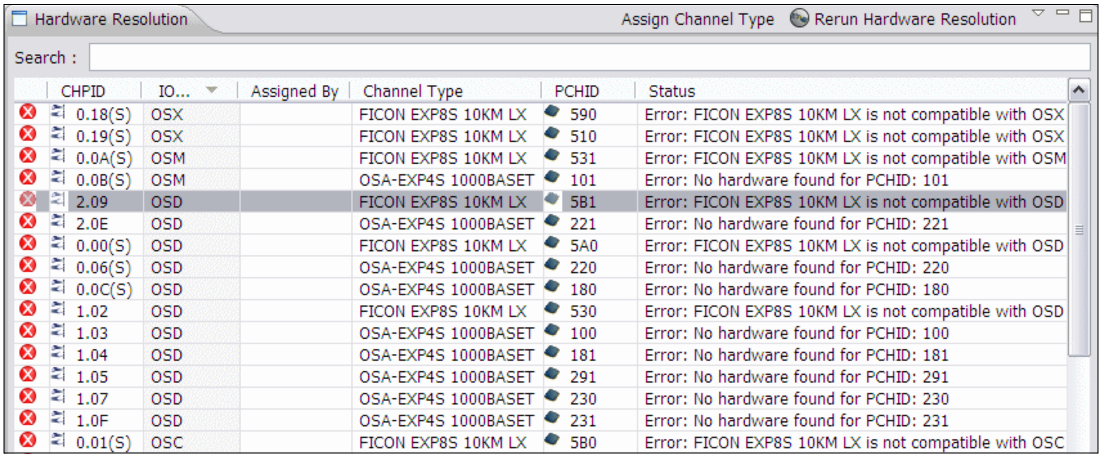

Figure 6-75 CMT - Channel_type is not compatible with IOCP_type

Select the channel type OSD. The Status is `Error: FICON EXP8S is not compatible with OSD`. Right-click in the row and select **Reset Incompatible (Hardware - I/O) Entries** to remove the PCHID values for only those rows (Figure 6-76).

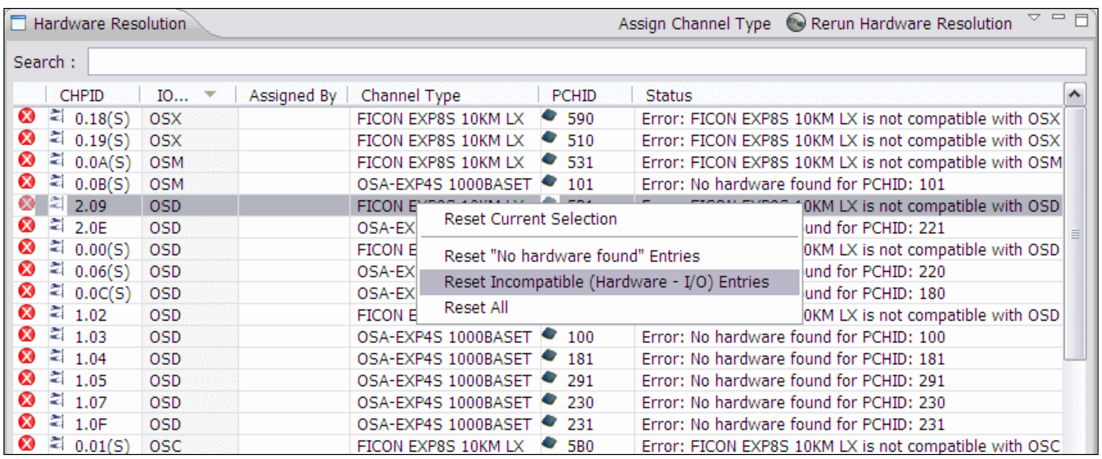

Figure 6-76 CMT - Channel_Type is not compatible with IOCP_type OSD

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 PCHIDs information (Figure 6-77).

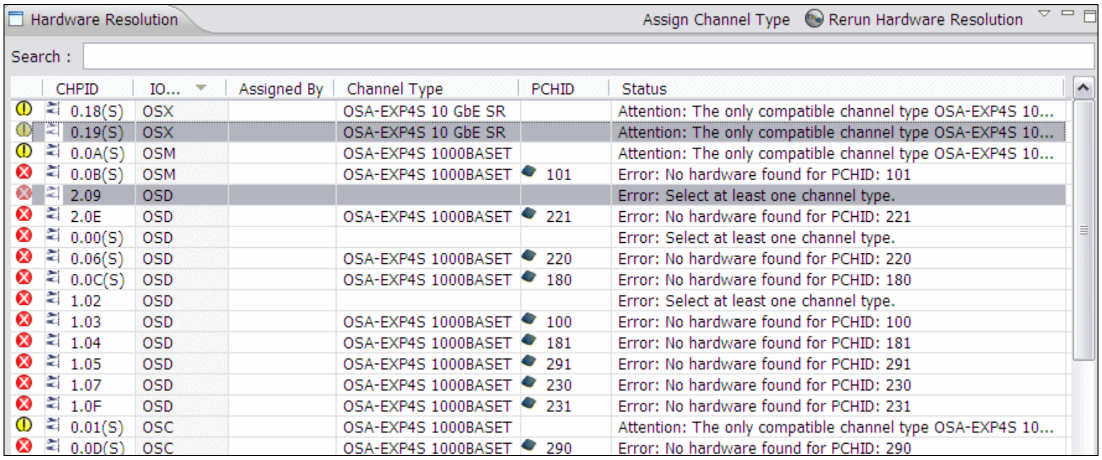

Figure 6-77 CMT - Results for reset of incompatible

The CHPID Mapping Tool now displays messages about any CHPID types that were imported from the IOCP input file (IODF) into the CMT that do not have any associated hardware support in the CFReport file (Figure 6-78). Click **OK**. The same figure also shows the Adapter Type Summary details.

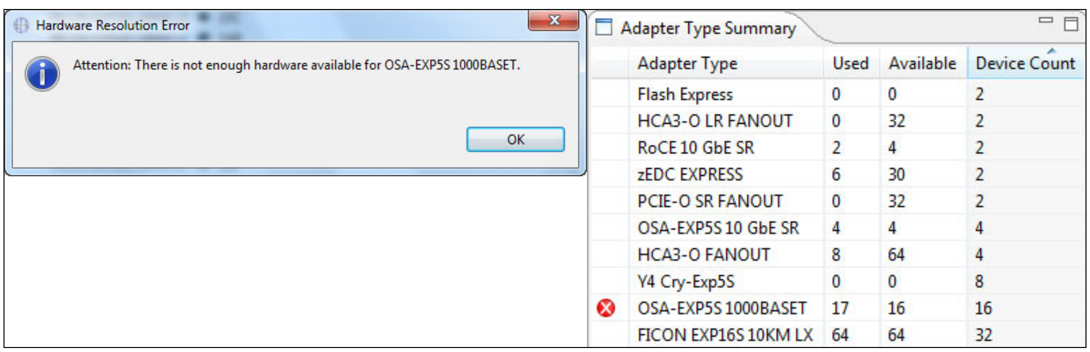

Figure 6-78 CMT - Required Hardware unavailable

Excessive numbers of *OSC* CHPID types are in the example IODF to show how the CHPID Mapping Tool handles this condition. For more information, see 6.2.8, “Overdefining channel paths on an XMP processor” on page 293.

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

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

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

To continue with this example, complete the following steps:

1. Return to the IODF and change the PCHID values for the OSD CHPIDs (or any other CHPIDs that have no 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 panel and selecting **Import IOCP File**.

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

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

Reset "Error: No hardware found" Entries

An X in a red circle in the first column indicates an error, and the Status column provides the information with the value of `Error: No hardware found` (Figure 6-79).

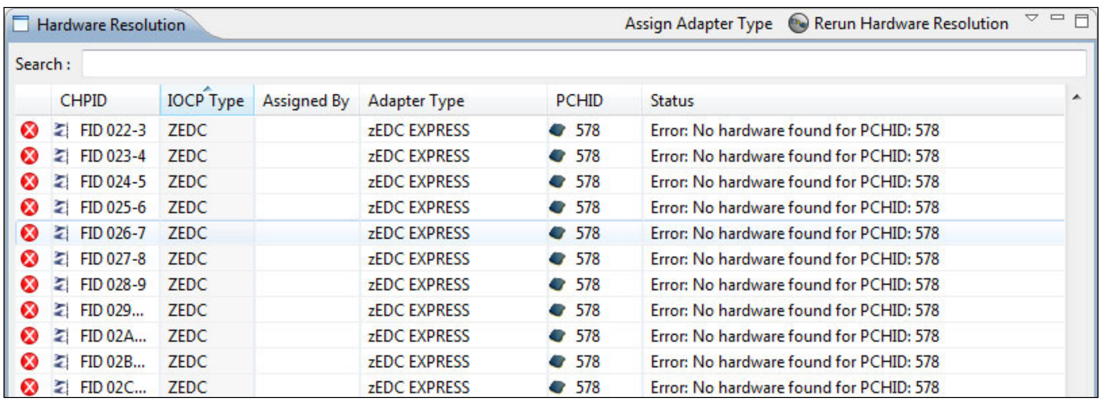

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 022-3	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 023-4	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 024-5	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 025-6	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 026-7	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 027-8	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 028-9	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 029...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02A...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02B...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578
FID 02C...	ZEDC		zEDC EXPRESS	578	Error: No hardware found for PCHID: 578

Figure 6-79 CMT - 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 PCHID values for those rows (Figure 6-80).

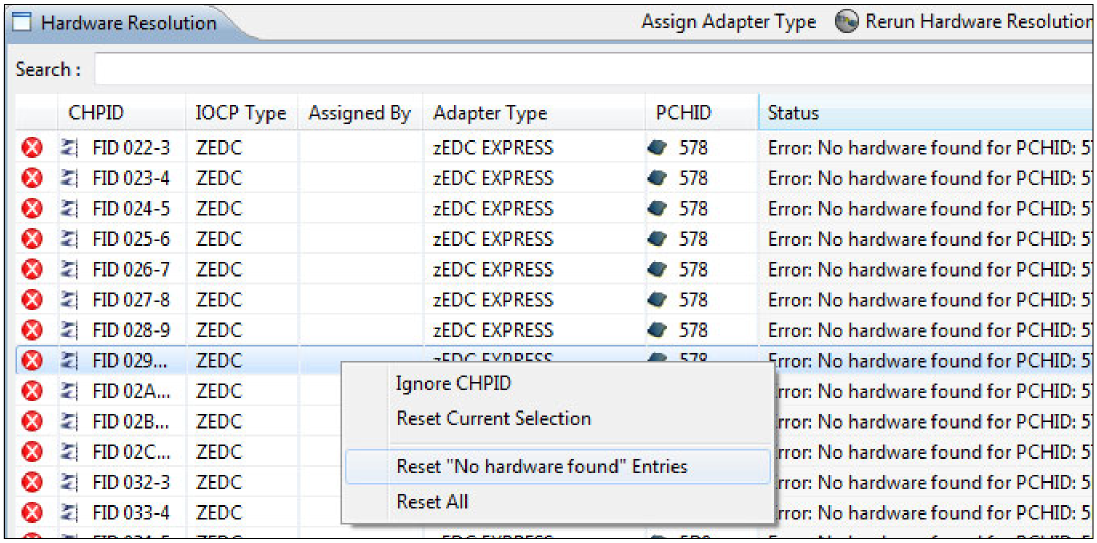

Figure 6-80 CMT - Resetting No Hardware found entries

The tool replaces the X with an *Attention* icon, changes the status message, and removes the PCHID information (Figure 6-81).

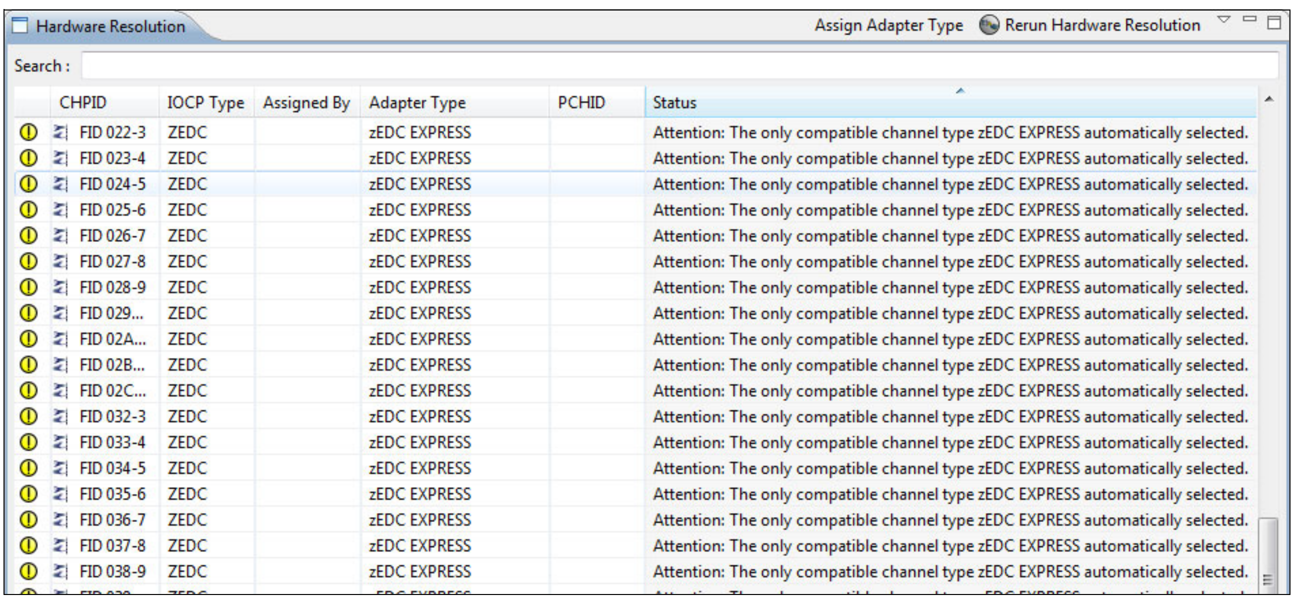

Figure 6-81 CMT - Results of resetting No hardware found

Reset “Select at least one adapter type”

The adapter type is not assigned to the current row. Assign an adapter type to the IOCP type:

1. Click the **Adapter Type** column in the target row. The tool displays an arrow in the Channel Type column of the target row (Figure 6-82).

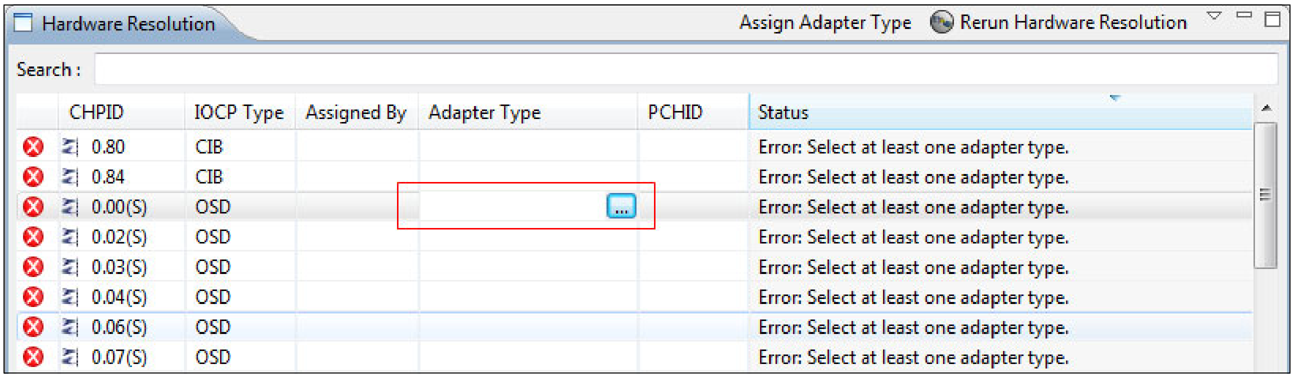

Figure 6-82 CMT - Selecting at least one adapter type

2. Click the ellipsis (...) box. The tool displays a list of available and compatible card types for the CHPID as shown in Figure 6-83.
3. Select an adapter type and click **OK**.
4. In the Adapter Type Summary tab, observe that the “Used” and “Available” totals change.

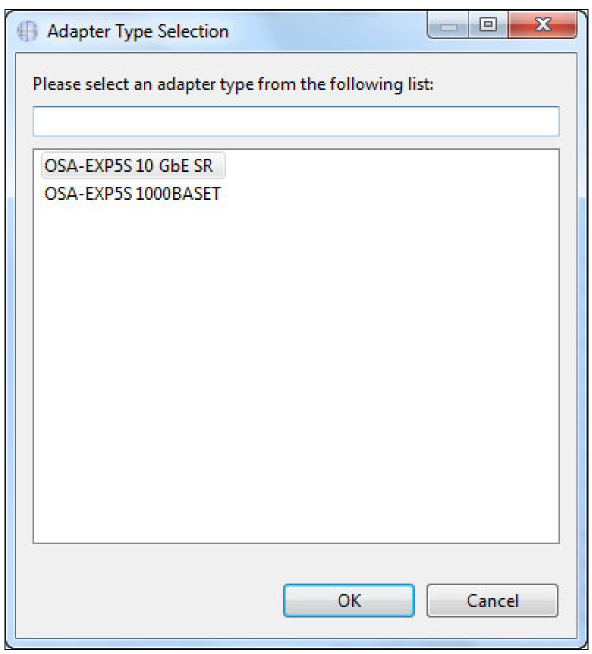

Figure 6-83 CMT - Adapter Type Selection

Reset “Required hardware for type IOCP_type not available”

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

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

You must change IOCP file or obtain more hardware.

Reset “PCHID_1 moved to new channel ID: PCHID_2”

When moving from old hardware to new hardware, for example during a miscellaneous equipment specification (MES), the PCHID value assigned to a feature may change. This message indicates that the IOCP file contains a PCHID value for the old machine that is being removed. The PCHID value is changed from the old machine to the PCHID value for the new machine. For example, PCHID_1 is the first PCHID value representing the old hardware (for example, 1B0) and PCHID_2 is the new value representing the new hardware (for example, 533). In essence, the feature is present in both the old and new hardware, but its location (PCHID) has 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** button becomes available.

6.6.5 Manual mapping to resolve CIB CHPIDs

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

To resolve the CIB CHPIDs, assign the available CHPIDs. Click **Manual Mapping** (Figure 6-84).

Figure 6-84 CMT - Manual Mapping

Ensure that the tool is set to display Manual Mapping in the **Hardware -> I/O** format (Figure 6-85).

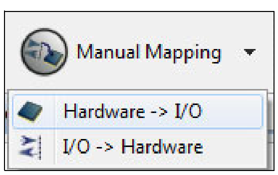

Figure 6-85 CMT - Manual Mapping of Hardware -> I/O

Click every row that has type HCA3-O in the Channel Type column. The tool displays all the available CHPIDs with IOCP type (Figure 6-86).

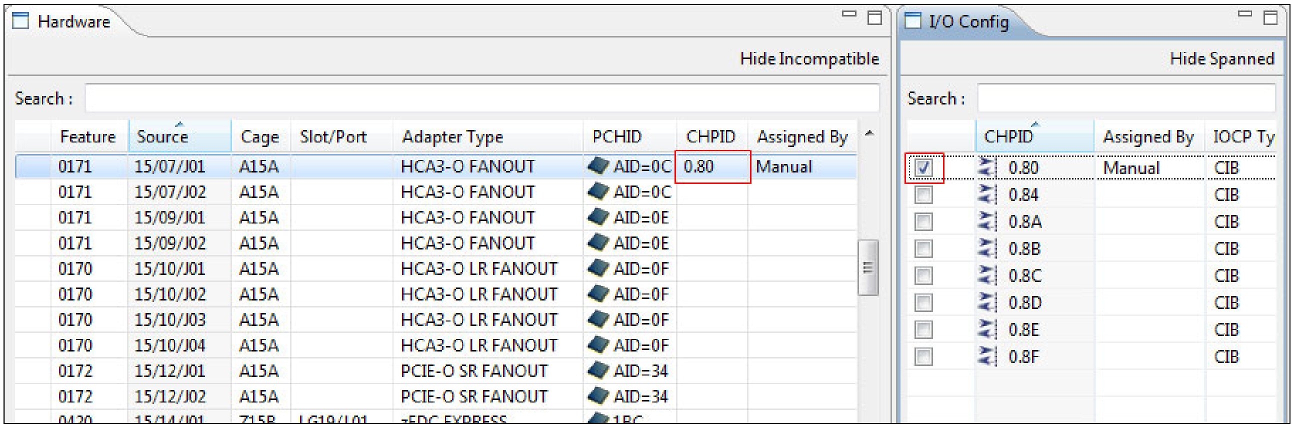

Figure 6-86 CMT - Adapter Type of HCA3 and associated CHPID assigned

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 HCA3-O adapter type, you see the `Multiple -->` value (Figure 6-87) inserted in the CHPID and Assigned By columns.

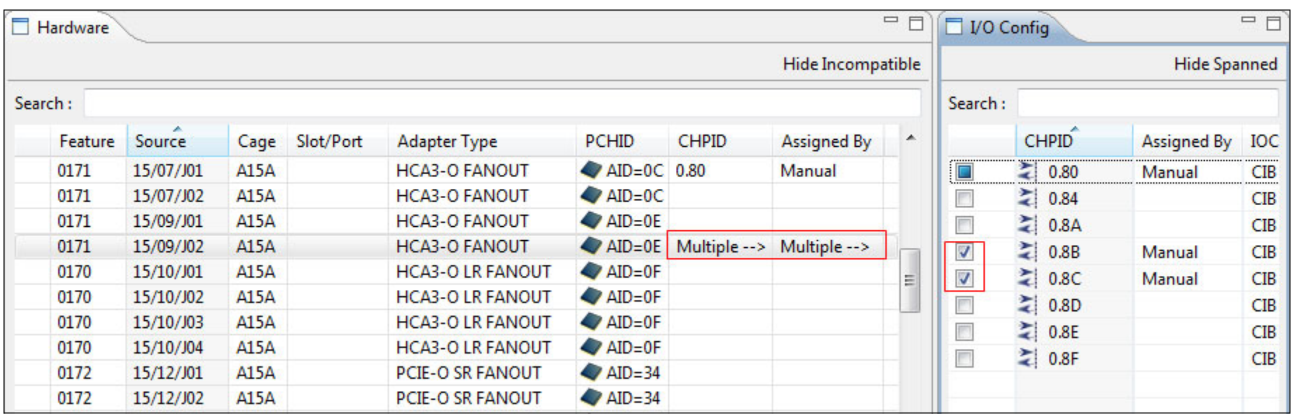

Figure 6-87 CMT - Adapter Type of HCA3 and associated multiple CHPID assigned

The **Automatic Mapping** button becomes available after you assign all the CHPIDs of IOCP type CIB.

6.6.6 Processing the CU Priority and Automatic Mapping

If you are importing an IOCP statements file from a 2817 or 2827 that had CU Priority values defined, review the CU Priority values first. The CHPID Mapping Tool can then perform the availability functions appropriately for a 2964.

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 these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices (Figure 6-88). For the example, select the following two options and then click **OK**:
   - **Reset CHPIDs assigned by Automatic Mapping**
   - **Reset CHPIDs assigned by IOCP**

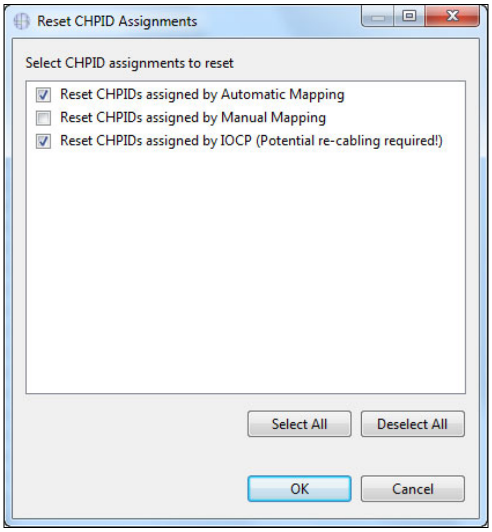

Figure 6-88 CMT - Reset CHPID Assignments

> **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 (Figure 6-89).

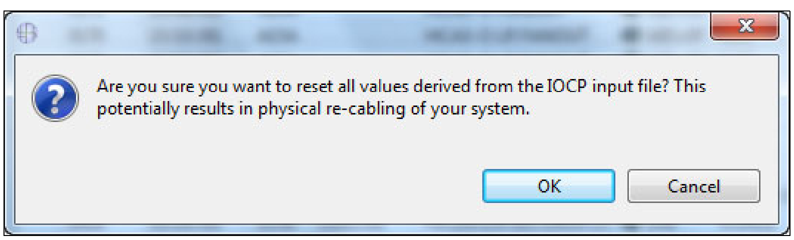

Figure 6-89 CMT - Reset CHPID assignments warning message

4. The 2964 has availability rules that differ from 2817 and 2827, so remove all PCHID 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 (Figure 6-90 on page 318). Click **OK**.

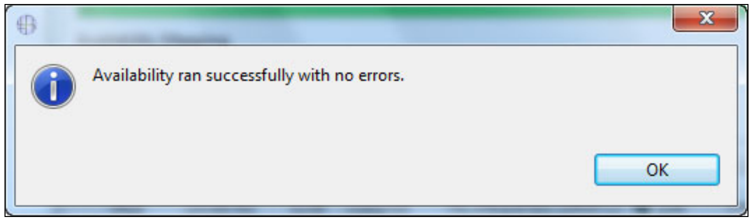

Figure 6-90 CMT - Availability ran successfully with no errors message

7. Click **OK** (Figure 6-91).

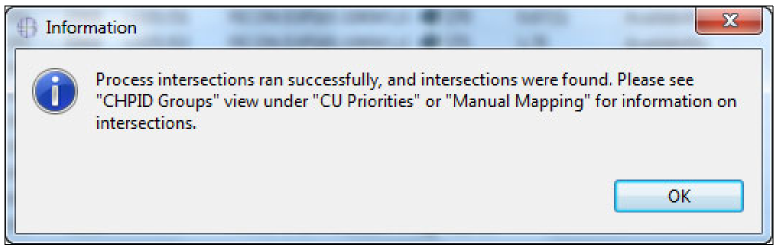

Figure 6-91 CMT - Process Intersections run successfully message

The possible intersects are as follows:

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 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 if they are acceptable (Figure 6-92). The example returned the “B” intersect.

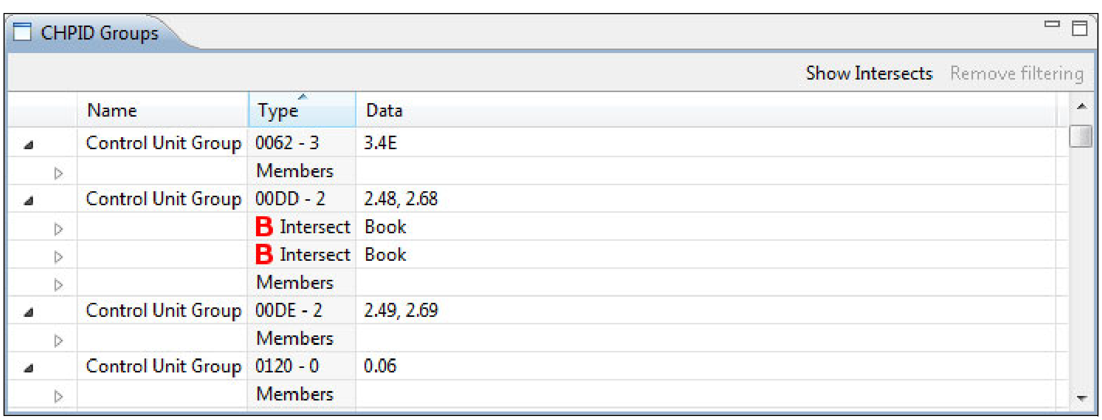

Figure 6-92 CMT - B Intersect examples

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.

Check and set values for items such as OSA-ICC CHPIDs and FCTC CHPIDs to ensure that the CHPID Mapping Tool allocates these CHPIDs with high PCHID availability.

1. Click **CU Priorities**. By default, this pane is in the center at the top.
2. In the CU Priorities pane, search in the CU Number column for the control units that you want to set a priority for.
3. Type 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. In the example, set the OSC type CU Numbers to priority 333 (Figure 6-93).

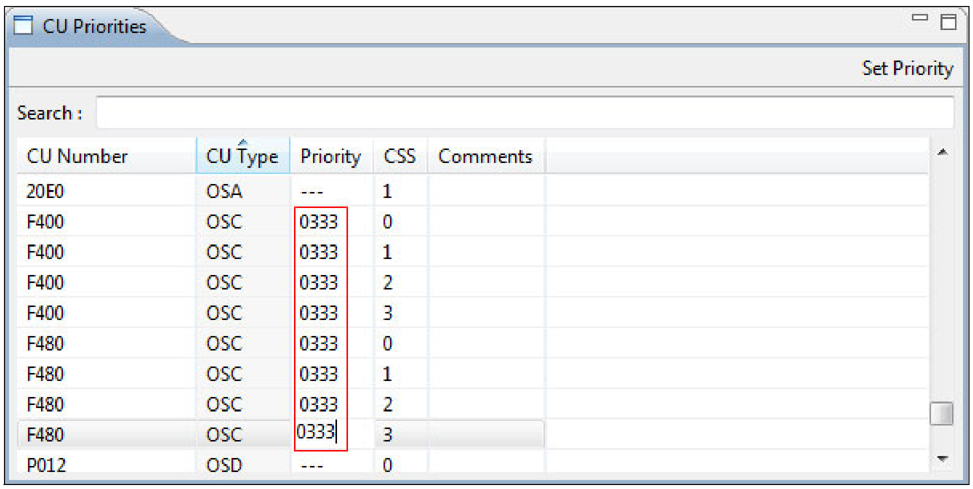

Figure 6-93 CMT - Set CU Priorities

If coupling links are used by a CF image, group those links.

Group each set of CHPIDs going to a different CPC with a common priority. For example, suppose the CF image has four links (CHPIDs 40, 41, 42, and 43) and that 40 and 41 go to one CPC, and 42 and 43 go to a different CPC. In this case, give CHPIDs 40 and 41 one priority and CHPIDs 42 and 43 a different priority. The concept is the same regardless of the number of connecting CPCs or the number of links to each CPC.

Now perform the second availability function by completing these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices. Click **Reset CHPIDs assigned by Automatic Mapping**.
3. Click **OK**.

   The Hardware Resolution pane shows that the CHPID and Assigned By columns are no longer blank (Figure 6-94 on page 320). The CMT assigned CHPIDs to PCHIDs and placed the Availability value in the Assigned By column, indicating that the CHPID values were assigned based on availability.

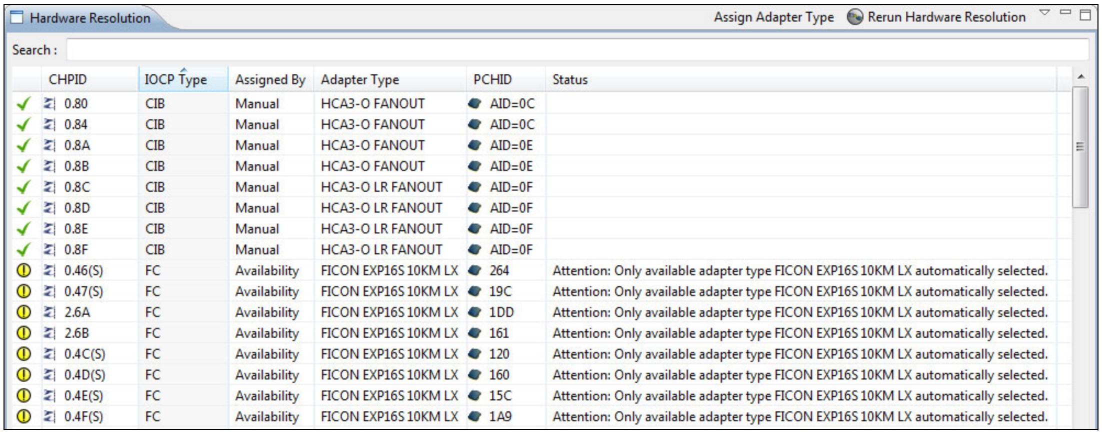

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
0.80	CIB	Manual	HCA3-O FANOUT	AID=0C	
0.84	CIB	Manual	HCA3-O FANOUT	AID=0C	
0.8A	CIB	Manual	HCA3-O FANOUT	AID=0E	
0.8B	CIB	Manual	HCA3-O FANOUT	AID=0E	
0.8C	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8D	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8E	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8F	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.46(S)	FC	Availability	FICON EXP16S 10KM LX	264	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.47(S)	FC	Availability	FICON EXP16S 10KM LX	19C	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
2.6A	FC	Availability	FICON EXP16S 10KM LX	1DD	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
2.6B	FC	Availability	FICON EXP16S 10KM LX	161	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4C(S)	FC	Availability	FICON EXP16S 10KM LX	120	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4D(S)	FC	Availability	FICON EXP16S 10KM LX	160	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4E(S)	FC	Availability	FICON EXP16S 10KM LX	15C	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4F(S)	FC	Availability	FICON EXP16S 10KM LX	1A9	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.

Figure 6-94 CMT - CHPIDs assigned

The possible Assigned By column values are as follows:

Manual You made the assignment by using manual mapping.

Automatic You made the assignment by using automatic mapping.

IOCP The IOCP source made the assignment.

Config File The CHPID Mapping Tool forced an assignment because of configuration file requirements.

> **More information:** See the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

4. You can now display the results of the channel mapping. You can also sort the report in various ways. For example, to see how the CHPID Mapping Tool ranked the control units, select the **CU Priorities** pane and click the **Priority** column (Figure 6-95).

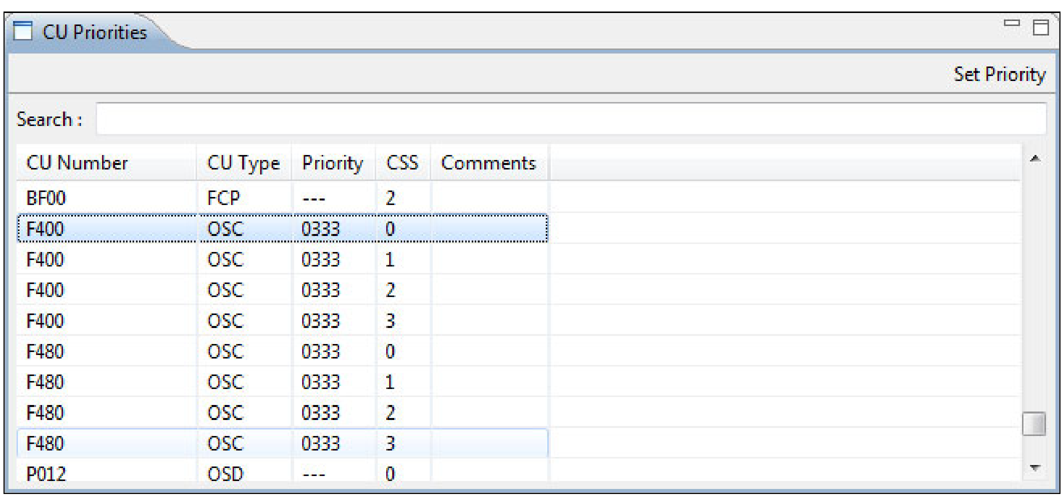

CU Number	CU Type	Priority	CSS	Comments
BF00	FCP	---	2	
F400	OSC	0333	0	
F400	OSC	0333	1	
F400	OSC	0333	2	
F400	OSC	0333	3	
F480	OSC	0333	0	
F480	OSC	0333	1	
F480	OSC	0333	2	
F480	OSC	0333	3	
P012	OSD	---	0	

Figure 6-95 CMT - CU Priorities showing assigned priorities

The example illustrates how CU Priority values are represented in the IOCP file.

Tip: The control unit priorities are stored in the IOCP output file created by the CMT that gets migrated back into HCD. HCD maintains these priorities and outputs them when it creates another IOCP deck. They are in the form of commented lines at the end of the IOCP deck, as shown here:

```
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
```

6.6.7 CHPIDs not connected to control units

In the CU Priorities window, click in the **CU Number** column (Figure 6-96). The CHPID Mapping Tool shows, at the end of the list, all CHPIDs 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.

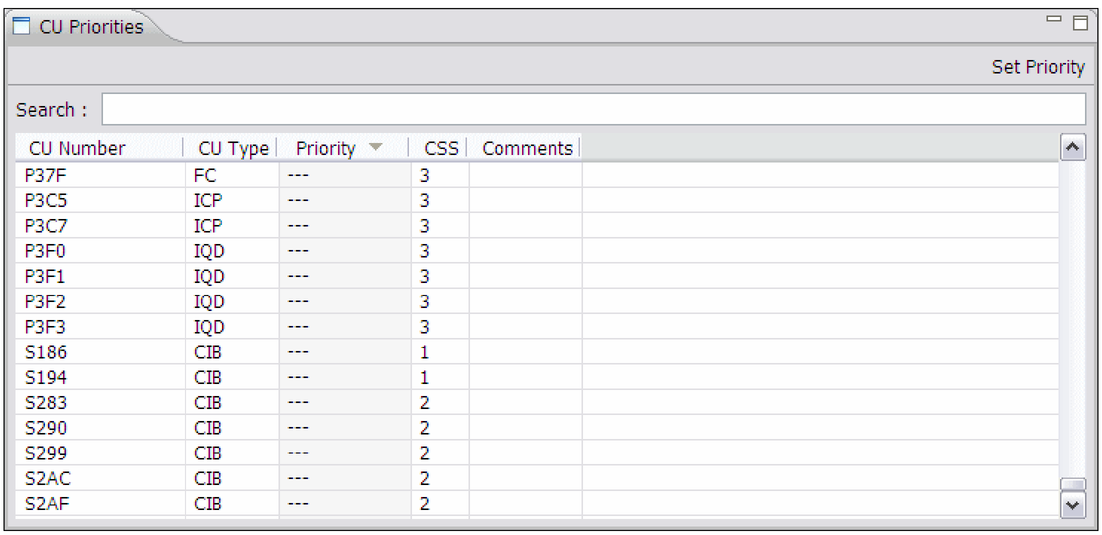

CU Number	CU Type	Priority	CSS	Comments
P37F	FC	---	3	
P3C5	ICP	---	3	
P3C7	ICP	---	3	
P3F0	IQD	---	3	
P3F1	IQD	---	3	
P3F2	IQD	---	3	
P3F3	IQD	---	3	
S186	CIB	---	1	
S194	CIB	---	1	
S283	CIB	---	2	
S290	CIB	---	2	
S299	CIB	---	2	
S2AC	CIB	---	2	
S2AF	CIB	---	2	

Figure 6-96 CMT - 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 need to 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 there are extra CHPIDs for anticipated new control units, consider grouping these CHPIDs with a common priority. Having a common priority allows the availability mapping function to pick PCHIDs that can afford your new control unit availability.

6.6.8 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**. Figure 6-97 shows the options to create a Preview Report or Save Report.

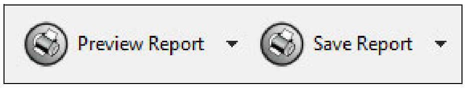

Figure 6-97 CMT - Preview Report and Save Report buttons

Click **Preview Report** or **Save Report** to display choices (a list of types of reports). The choices are the same except Save Report lists an extra selection (Figure 6-98).

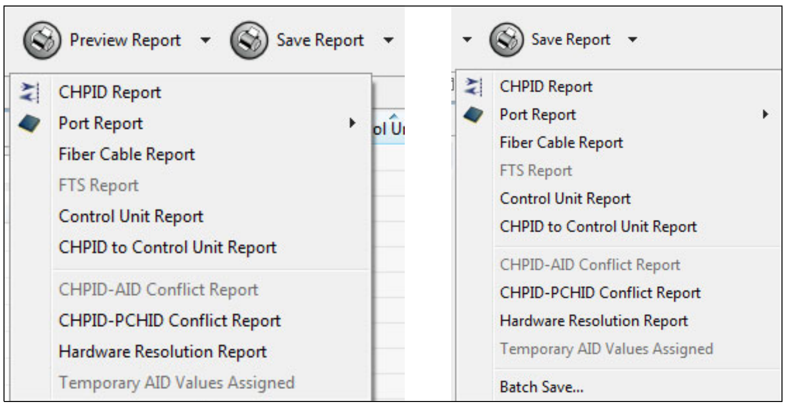

Figure 6-98 CMT - Preview Report and Save Report menus

For simplicity, only three reports are described in this example: The CHPID Report; the Port Report, sorted by location; and the 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 PCHID or cage/slot/port information before system delivery.

CHPID Report

To create the CHPID report, complete the following steps:

1. Click **Preview Report** → **CHPID Report** (Figure 6-99).

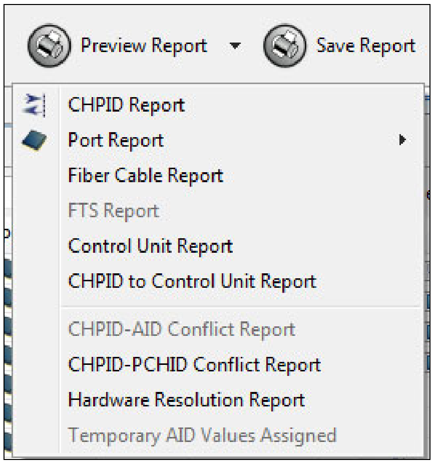

Figure 6-99 CMT - Preview report - CHPID Report

The CHPID Mapping Tool displays the CHPID Report in a Report tab within the CMT (Figure 6-100).

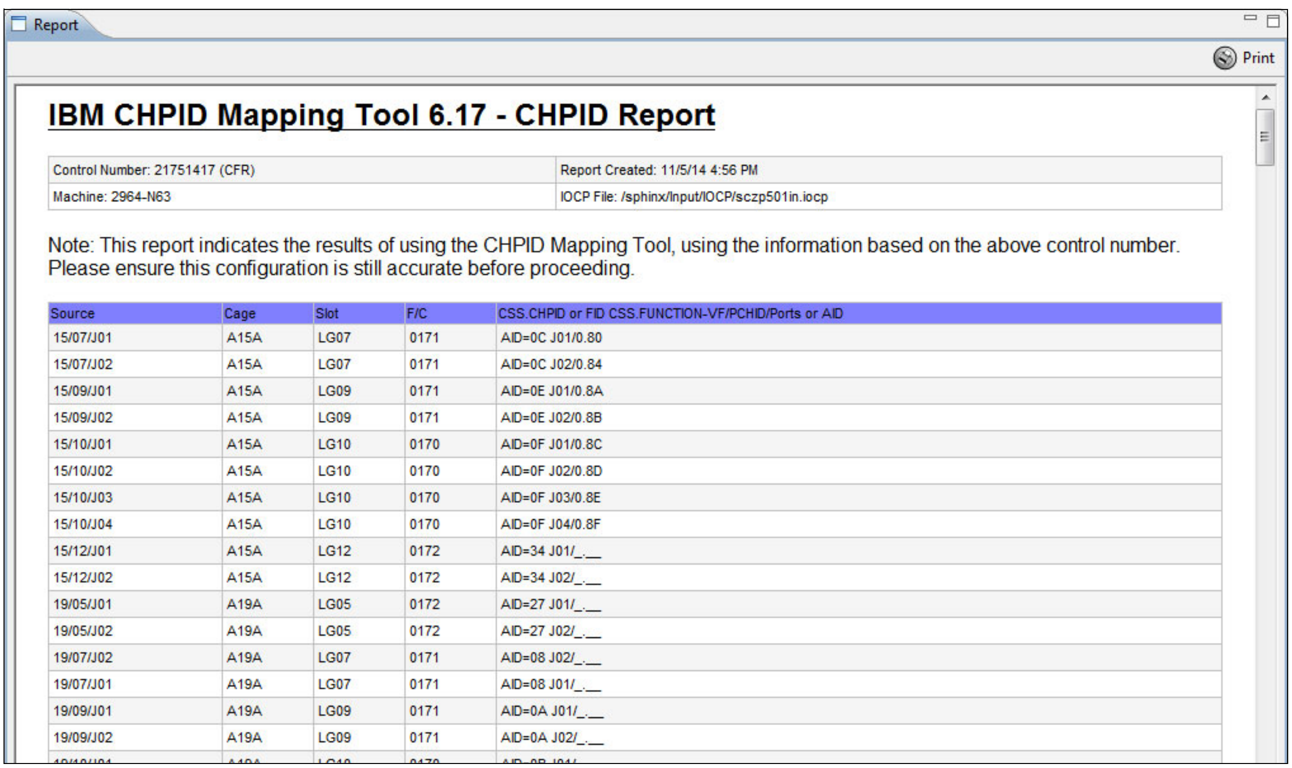

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__

Figure 6-100 CMT - 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 (Figure 6-101). 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**.

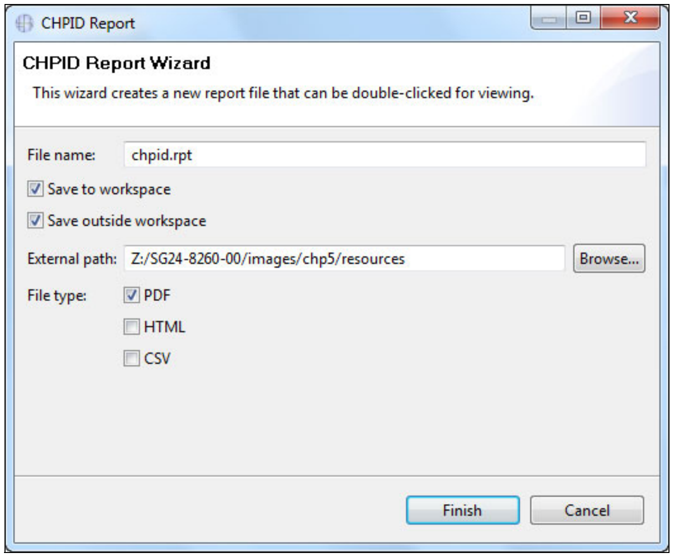

Figure 6-101 CMT - Save CHPID Report

The CHPID Report is created by the CHPID Mapping Tool (Figure 6-102).

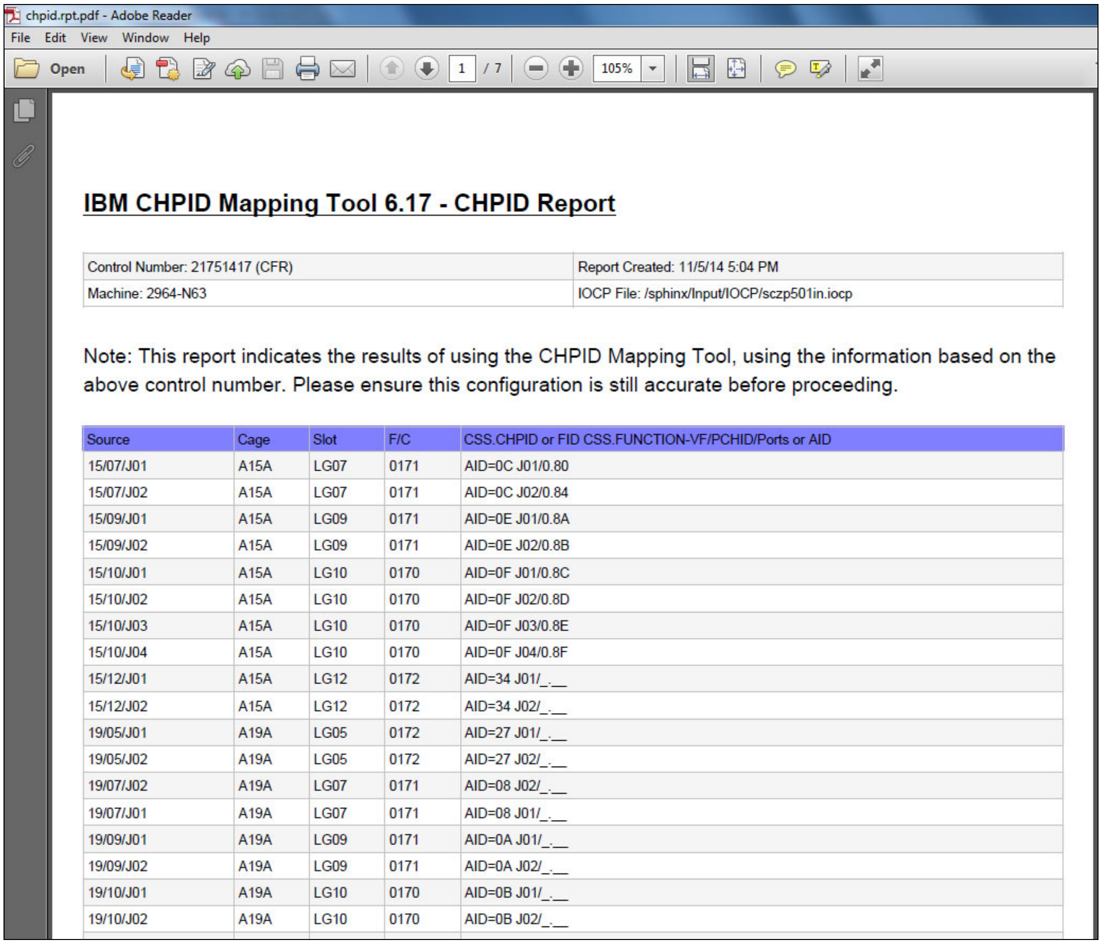

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:04 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__
19/10/J01	A19A	LG10	0170	AID=0B J01/_.__
19/10/J02	A19A	LG10	0170	AID=0B J02/_.__

Figure 6-102 CMT - CHPID Report example in PDF format

At the end of this CHPID Report is a list of CHPIDs with modified PCHID/AID assignments (Figure 6-103). This report is valuable for moving cables.

List of CHPIDs having modified PCHID/AID assignments

Note: For CHPIDs that had PCHID/AID assignments in the IOCP file that was loaded for this session of the Mapping Tool.

CHPIDs or FUNCTIONs	Previous PCHID/AID-Port	PCHID/AID-Port	Current Location	F/C
0.80	09-1	0C-1	A15ALG07J.01	0171
0.84	09-1	0C-02	A15ALG07J.02	0171
0.8A	09-1	0E-1	A15ALG09J.01	0171
0.8B	09-1	0E-02	A15ALG09J.02	0171
0.8C	09-1	0F-01	A15ALG10J.01	0170
0.8D	09-1	0F-02	A15ALG10J.02	0170
0.8E	09-1	0F-03	A15ALG10J.03	0170
0.8F	09-1	0F-04	A15ALG10J.04	0170
FID 000	544	208	Z08BLG03J.01	0411
FID 010	5EC	140	Z22BLG20J.01	0411
FID 020-1	578	1BC	Z15BLG19J.01	0420
FID 021-2	578	27C	Z08BLG38J.01	0420
FID 022-3	578	Not Assigned		
FID 023-4	578	Not Assigned		

Figure 6-103 CMT - List of CHPIDs that have modified PCHID/AID assignments

CHPID to Port Report, sorted by location

To create the Port Report, 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 (Figure 6-104).

IBM CHPID Mapping Tool 6.17 - CHPID to Port Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:12 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
A15A	07/LG	AID=0C / J.01	15/07	HCA3-O FANOUT	0.80	Manual
A15A	07/LG	AID=0C / J.02	15/07	HCA3-O FANOUT	0.84	Manual
A15A	09/LG	AID=0E / J.01	15/09	HCA3-O FANOUT	0.8A	Manual
A15A	09/LG	AID=0E / J.02	15/09	HCA3-O FANOUT	0.8B	Manual
A15A	10/LG	AID=0F / J.01	15/10	HCA3-O LR FANOUT	0.8C	Manual
A15A	10/LG	AID=0F / J.02	15/10	HCA3-O LR FANOUT	0.8D	Manual
A15A	10/LG	AID=0F / J.03	15/10	HCA3-O LR FANOUT	0.8E	Manual
A15A	10/LG	AID=0F / J.04	15/10	HCA3-O LR FANOUT	0.8F	Manual
A15A	12/LG	AID=34 / J.01	15/12	PCIE-O SR FANOUT		
A15A	12/LG	AID=34 / J.02	15/12	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.01	19/05	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.02	19/05	PCIE-O SR FANOUT		
A19A	07/LG	AID=08 / J.01	19/07	HCA3-O FANOUT		
A19A	07/LG	AID=08 / J.02	19/07	HCA3-O FANOUT		
A19A	09/LG	AID=0A / J.01	19/09	HCA3-O FANOUT		

Figure 6-104 CMT - CHPID to Port Report, sorted by location

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 (Figure 6-105).

IBM CHPID Mapping Tool 6.17 - CHPID to CU Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:15 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
0	00	OSD	19/12/J01	Z08B LG07 J00J01	214	2040	OSA	---
0	01	OSC	19/03/J01	Z15B LG21 J00J01	1C4	F400	OSC	0333
0	02	OSD	15/03/J01	Z15B LG35 J00J01	1F0	2080	OSA	---
0	03	OSD	15/15/J01	Z22B LG17 J00J01	134	20A0	OSA	---
0	04	OSD	15/02/J01	Z22B LG36 J00J01	174	20C0	OSA	---
0	06	OSD	19/02/J01	Z22B LG21 J00J01	144	0120	OSA	---
0	07	OSD	15/02/J01	Z22B LG37 J00J01	178	0140	OSA	---
0	0C	OSD	19/03/J01	Z15B LG22 J00J01	1C8	2060	OSA	---
0	0D	OSC	15/15/J01	Z22B LG18 J00J01	138	F480	OSC	0333
0	12	OSD	15/05/J01	Z08B LG36 J00J01	274	P012	OSD	---
0	13	OSD	19/14/J01	Z15B LG06 J00J01	190	P013	OSD	---
0	40	FC	15/14/J01	Z15B LG12 D1	1A4	1000	2107	---
						1200	2107	---
						1400	2107	---

Figure 6-105 CMT - CHPID to CU Report

6.6.9 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 has PCHIDs that are assigned to CHPIDs.

Using HCM: You might prefer to use HCM to transfer the updated IOCP statements file back to the host. However, first run the next step in the CHPID Mapping Tool to create the updated IOCP file.

To create the IOCP, complete the following steps:

1. Select **File** → **Export IOCP input file** (Figure 6-106).

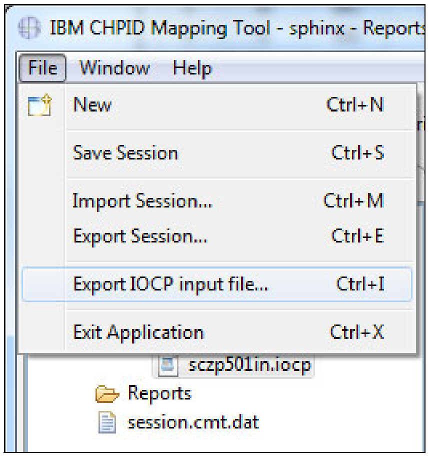

Figure 6-106 CMT - Export IOCP input file

2. Enter the Export Path and IOCP Name for the IOCP output file and click **Finish** (Figure 6-107).

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

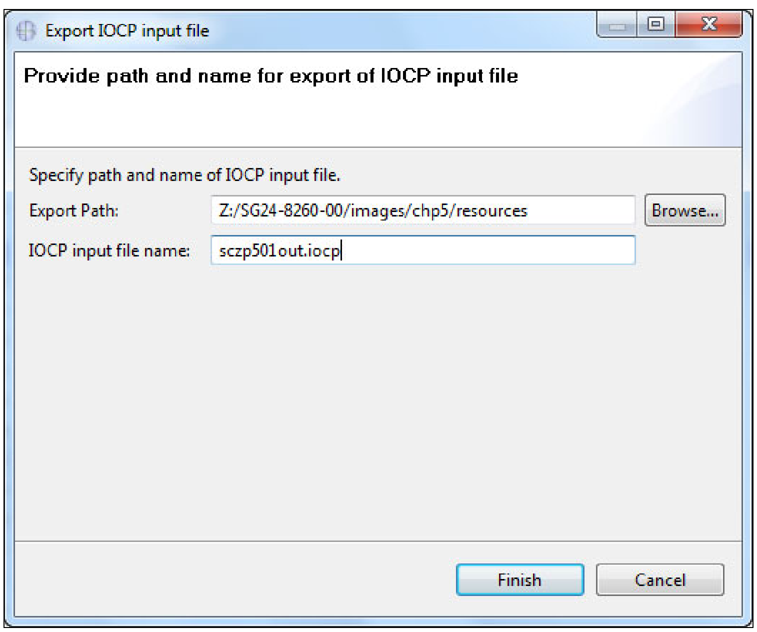

Figure 6-107 CMT - Export IOCP File

3. Select **File** → **Save Session** (Figure 6-108).

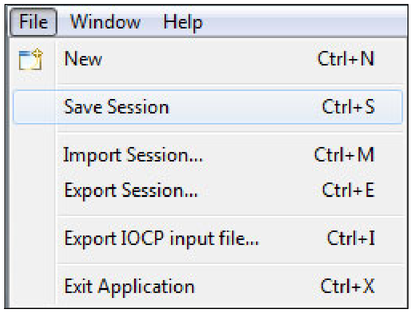

Figure 6-108 CMT - Save session

You might want to save your project before exiting the CMT application.

6.7 HCD: Updating the 2964 work IODF with PCHIDs

After you map the PCHIDs to CHPIDs by using the CHPID Mapping Tool, transfer this information back into HCD. To update the IODF with the PCHIDs, follow these following steps:

1. Upload the IOCP file that was created by the CMT (`sczp501out.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 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. They are in the form of commented lines at the end of the IOCP deck (Example 6-6).

Example 6-6 HCD - Updated IOCP statements file: with CMT statements

```
         CNTLUNIT CUNUMBR=FFFD,PATH=((CSS(0),C6,C7)),UNIT=CFP
         IODEVICE ADDRESS=(FFCF,007),CUNUMBR=(FFFD),UNIT=CFP
         IODEVICE ADDRESS=(FFDD,007),CUNUMBR=(FFFD),UNIT=CFP
         CNTLUNIT CUNUMBR=FFFE,PATH=((CSS(0),C4,C5)),UNIT=CFP
         IODEVICE ADDRESS=(FFD6,007),CUNUMBR=(FFFE),UNIT=CFP
         IODEVICE ADDRESS=(FFF2,007),CUNUMBR=(FFFE),UNIT=CFP
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
******************************** Bottom of Data ************************
```

2. From the HCD main panel, enter the work IODF name used. Select option **5. Migrate configuration data** (Figure 6-109).

```
                             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 . . . 'SYS6.IODFB6.WORK'                     +
```

Figure 6-109 HCD - Hardware Configuration: migrate configuration data.

3. From the Migrate Configuration Data panel (Figure 6-110), 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
```

Figure 6-110 HCD - Migrate Configuration Data: migrate IOCP/OS data.

4. The Migrate IOCP Data panel opens (Figure 6-111 on page 331). 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 OS configuration that is associated with the processor.

IOCP only input data set This data set was specified when the `iocpout.txt` file was uploaded to z/OS.

Processing mode Select option **2** to save the results of the migration. (Before using option 2, however, 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 . . . . . . . . . . . . SCZP501_  +   CSS ID . . . . . .    +
OS configuration ID  . . . . . . . . L06RMVS1  +

Combined IOCP/MVSCP input data set . _____________________________________
IOCP only input data set . . . . . . 'SYS6.IODFB6.IOCPOUT.SCZP501'
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)
```

Figure 6-111 HCD - 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 (Figure 6-112).

```
------------------------- 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
#                     SYS6.IODFB6.WORK.
**************************** Bottom of data *****************************
```

Figure 6-112 HCD - Migration Message List: successful message

The work IODF now contains both the CHPID definitions and the mapping to PCHIDs 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.

6.8 HCD: Building the 2964 production IODF

To use the definitions that were updated in HCD, create a 2964 production IODF from your work IODF. Then, remotely or locally write the IODF to the 2964 IOCDS by using Write IOCDS in preparation for the upgrade.

Complete the following steps:

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 6-113).

```
                            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.IODFB6.WORK'                     +
```

Figure 6-113 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 6-114). 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
```

Figure 6-114 HCD - Activate or Process Configuration Data: build production IODF

3. HCD displays the Message List panel (Figure 6-115). Verify that you have only severity “W” (warning) messages and that they are normal for your configuration. Correct any other messages that should not occur and try to build the production IODF again. Continue this process until you have no messages that indicate problems.

```
------------------------------ Message List -------------------------------
  Save  Query  Help
--------------------------------------------------------------------------
                                                                Row 1 of 18
Command ===> ___________________________________________ Scroll ===> CSR

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

/ Sev Msg. ID  Message Text
_ W   CBDA857I No channel paths attached to partition A0F of processor
#              SCZP501.0.
_ W   CBDA857I No channel paths attached to partition A1B of processor
#              SCZP501.1.
_ W   CBDA857I No channel paths attached to partition A1D of processor
#              SCZP501.1.
_ W   CBDA857I No channel paths attached to partition A1E of processor
#              SCZP501.1.
_ W   CBDA857I No channel paths attached to partition A1F of processor
#              SCZP501.1.
_ W   CBDA857I No channel paths attached to partition A3B of processor
```

Figure 6-115 HCD - Message List: building production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (Figure 6-116). Complete the Production IODF name and Volume serial number fields, and then press Enter.

```
------------- Build Production I/O Definition File ------


Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODFB6.WORK'

Production IODF name . 'SYS6.IODFB6'
Volume serial number . IODFPK  +

Continue using as current IODF:
2   1.  The work IODF in use at present
    2.  The new production IODF specified above
```

Figure 6-116 HCD - Build Production I/O Definition File: data fields to be updated

6. The Define Descriptor Fields panel opens (Figure 6-117). 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.IODFB6'

Descriptor field 1  . . . SYS6
Descriptor field 2  . . . IODFB6
```

Figure 6-117 HCD - 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.IODFB6`** `created.`

6.9 HCD/HMC: Loading the 2964 processor IOCDS

You now have a production IODF, named `SYS6.IODFB6`. Now the IOCDS component of the IODF must be updated on the replacement CPC that is being installed (for example, SCZP501) and activated (POR) using this IOCDS.

The final step is to perform an initial program load (IPL) of the processor using this IODF. (Describing how to perform the IPL of the new hardware is beyond the scope of this book.)

The two possible ways to load the IOCP Statements onto the 2964 Support Element IOCDS are as follows:

- HCD, by using option 2.11
- The HMC/SE, by using the Stand-Alone Input/Output Configuration Program

Although both are valid methods to write the new configuration to the IOCDS, using HCD option 2.11 is the preferred method. However, your 2964 processor and Support Element that are replacing the 2827 might not be connected to the system where the configuration was generated or cannot be connected to any system where HCD is running. In that case, use the Stand-Alone IOCP process.

6.9.1 Updating the IOCDS using HCD option 2.11

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

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 6-118). Ensure that the IODF is the production IODF that was created in 6.8, “HCD: Building the 2964 production IODF” on page 332 and then press Enter.

```
                            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.IODFB6'                        +
```

Figure 6-118 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 6-119). Select option **11. Build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes**.

```
----- 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
```

Figure 6-119 HCD - Activate or Process Configuration data: build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes

Consideration: In this example, the replacement 2964 has been delivered and has connectivity to the HMC LAN so you can create an IOCDS from which to power-on reset. This might not be the case for all situations.

If the replacement 2964 is not accessible from the HMC LAN, copy the IOCP statements onto a USB flash drive and import them into the 2964 HMC to run a stand-alone IOCP. You can create a file on a USB flash drive by using the same process that is used to create an IOCP input file for the CHPID Mapping Tool.

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

3. The System z Cluster List panel opens (Figure 6-120). In the list, use a forward slash (/) to select the new 2964 to update one of its IOCDSs. Then, press Enter.

```
                            System z Cluster List                   Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address       Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
# USIBMSC.SCZP401    2827   H43
/ USIBMSC.SCZP501    2964   N63     SCZP501
******************************* Bottom of data ********************************
```

Figure 6-120 HCD - System z Cluster List: selecting processor for IOCDS update

4. The Actions on selected CPCs panel opens (Figure 6-121). 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)
```

Figure 6-121 HCD - Actions on selected CPCs: work with IOCDSs

5. The IOCDS List panel opens (Figure 6-122). Select the IOCDS that you want to update for the 2964 installation by typing a forward slash (/) next to it, and 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.SCZP501  IODFB5   LPAR   Alternate  No        No          No
_ A1.SCZP501  IODF40   LPAR   Alternate  No        No          No
_ A2.SCZP501  IODF42   LPAR   Alternate  No        No          No
_ A3.SCZP501           LPAR   POR        Yes       No          Yes-POR
****************************** Bottom of data *******************************
```

Figure 6-122 HCD - IOCDS List: selecting IOCDS for update

6. The Actions on selected IOCDSs panel opens (Figure 6-123). Select option **1. Update IOCDS** and press Enter.

```
--------------- Actions on selected IOCDSs ------


Select by number or action code and press Enter.

1_  1.  Update IOCDS . . . . . . . . . . . . (u)
    2.  Switch IOCDS . . . . . . . . . . . . (s)
    3.  Enable write protection  . . . . . . (e)
    4.  Disable write protection . . . . . . (w)
```

Figure 6-123 HCD - Actions on selected IOCDSs: update IOCDS

7. The Build IOCDSs panel opens (Figure 6-124). Verify that all the information is correct. Complete the Title1 field and press Enter.

```
------------------------------ Build IOCDSs -------------------------------
                                                                 Row 1 of 1
Command ===> ___________________________________________ Scroll ===> CSR

Specify or revise the following values.

IODF name . . . . . . . . . : 'SYS6.IODFB6'

Title1 . IODFB6__________________________________________________________
Title2 : SYS6.IODFB6 - 2014-11-10 11:37

                           Write IOCDS in
IOCDS        Switch IOCDS  preparation of upgrade
A0.SCZP501   No            No
**************************** Bottom of data *****************************
```

Figure 6-124 HCD - Build IOCDSs: verifying IODF

8. The Job Statement Information panel opens (Figure 6-125 on page 338). Complete the job statements as required by the installation and press Enter. HCD submits the job to update the IOCDS.

```
----------------------- Job Statement Information ---------------


Specify or revise the job statement information.


Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 6-125 HCD - Job Statement Information: option to override job statement cards

9. Verify the job output to ensure that the IOCDS was written without error and to the correct IOCDS. You receive a message similar to the following message:

```
ICP057I IOCP JOB WIOCP    SUCCESSFUL.  LEVEL A0 IOCDS REPLACED.

Sev  Msgid    Message Text
 I  CBDA674I  IOCP successfully completed for A0.SCZP501.
```

10. Now if you return to HCD option 2.11 and view the IOCDS, the SNA Address is at `USIBMSC.SCZP501` (Figure 6-126).

```
                            System z Cluster List                  Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
# USIBMSC.SCZP401    2827   H43
/ USIBMSC.SCZP501    2964   N63     SCZP501
******************************* Bottom of data ********************************
```

Figure 6-126 HCD - System z Cluster List: selecting processor for IOCDS verify

Figure 6-127 shows the updated IOCDS with Alternate status.

```
                                  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.SCZP501  IODFB6   LPAR   Alternate  No        Yes         No
_ A1.SCZP501  IODF40   LPAR   Alternate  No        No          No
_ A2.SCZP501  IODF42   LPAR   Alternate  No        No          No
_ A3.SCZP501  IODF43   LPAR   POR        Yes       No          Yes-POR
***************************** Bottom of data ********************************
```

Figure 6-127 HCD - IOCDS List: IOCDS verified

6.9.2 Updating the IOCDS using Stand-Alone Input/Output Config Program

Copy the IOCP statements that were generated by using HCD option **2.3. Build IOCP input data set** onto a USB flash memory drive and retain it.

Tip: For more information, see the *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152.

To update the IOCDS, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation, supplied with the 2964, as opposed to a remote web browser. Select the new 2964, assuming it was defined to the Defined CPCs Work Area.
2. Perform a power-on reset using one of the Reset profiles and Starter IOCDSs provided with the processor during installation. This action creates an environment on the processor in which you can run the Stand-Alone IOCP process.
3. When the power-on reset is complete, activate one of the logical partitions with at least 128 MB of storage. Use this partition to run the I/O Configuration Program.
4. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
5. Under Systems or Members, click the system to select it (in this example, SCZP501).
6. On the Tasks tab, click **Recovery** → **Single Object Operations** → **Yes**.
7. Under Systems Management, click the system to select it (in this example, SCZP501).

8. Under Partitions, select the LPAR you want to use to run the Stand-Alone IOCP program (in this example, A0B) as shown in Figure 6-128.

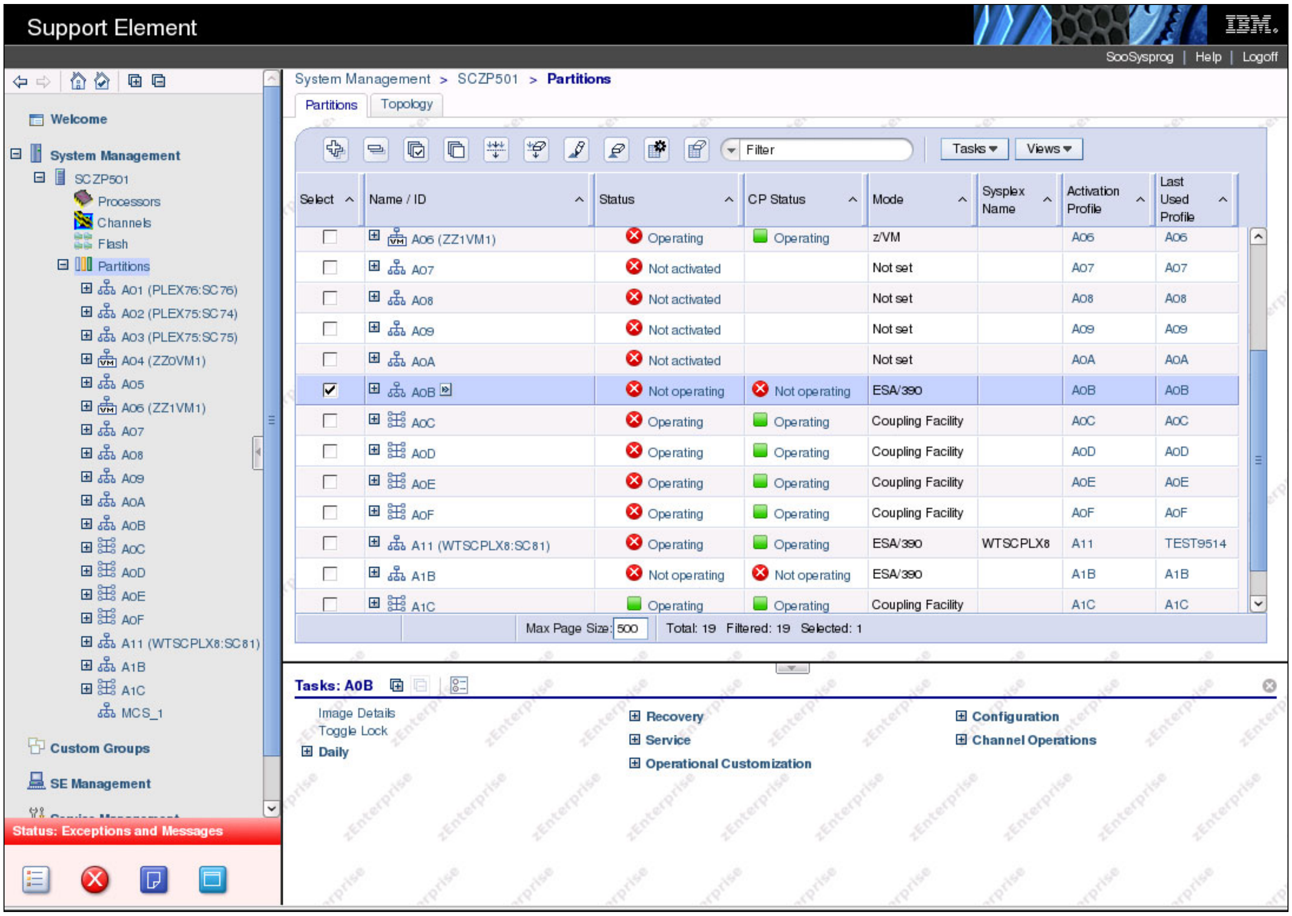

Figure 6-128 HMC - SAIOCP: partition selected for SAIOCP program load

9. On the Tasks tab, click **Configuration** → **Input/output (I/O) Configuration**.

10. Select the data set into which you want to import the IOCDS (in this example, A0) as shown in Figure 6-129.

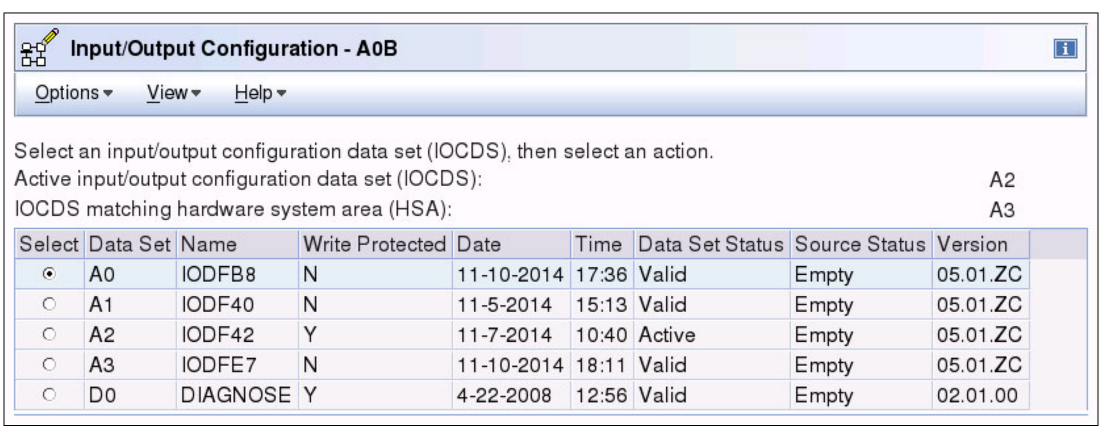

Figure 6-129 HMC - SAIOCP: IOCDS selection for import

11. Insert the USB flash memory drive that contains the IOCP text file. Wait for the drive insertion message to open (Figure 6-130).

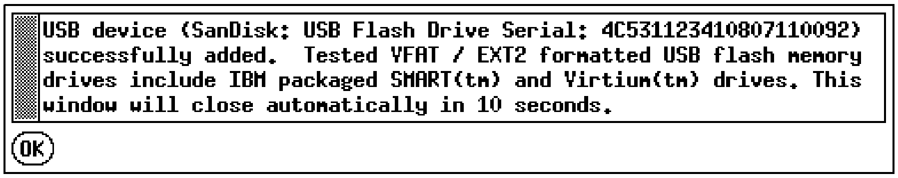

Figure 6-130 HMC - SAIOCP: USB drive insertion message

Tip: Only files in the root directory of the USB drive can be read by the HMC. Any folders and their contents will not be read.

12. Click **Options** → **Import Source File** → **Hardware Management Console USB Flash Memory Drive** (Figure 6-131).

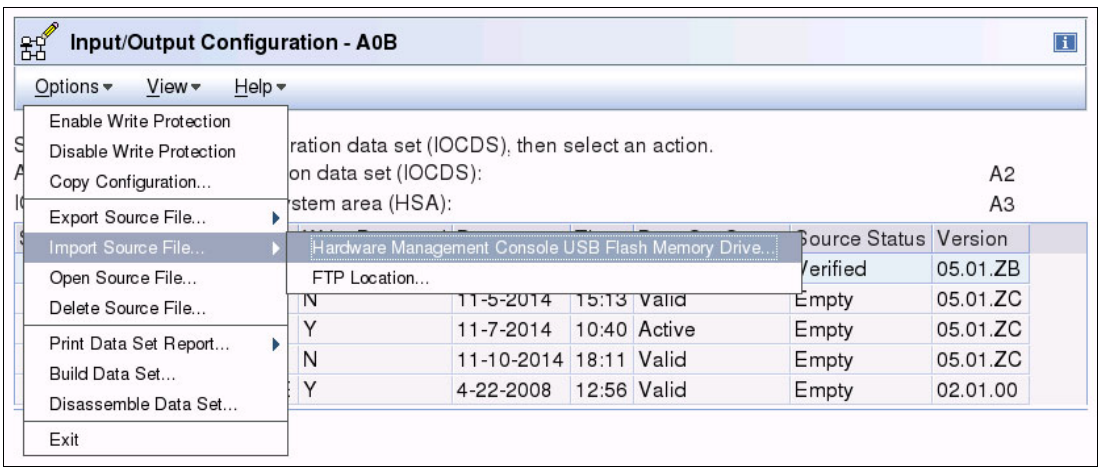

Figure 6-131 HMC - SAIOCP: import source file

13. Select the source file name and click **OK** (Figure 6-132).

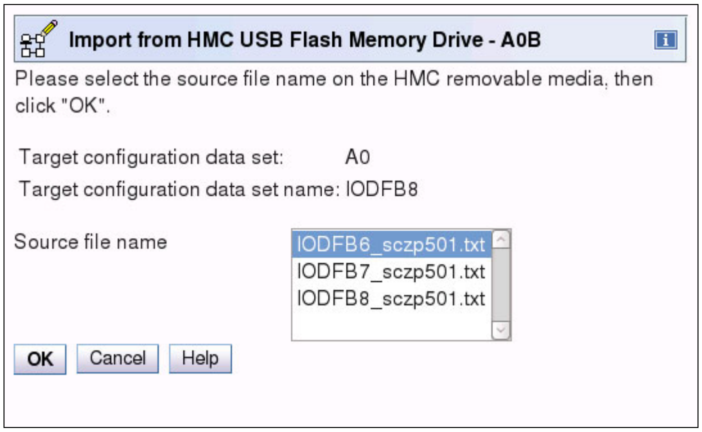

Figure 6-132 HMC - SAIOCP: select source file

14. The source file is now read from the USB drive. Click **OK** (Figure 6-133).

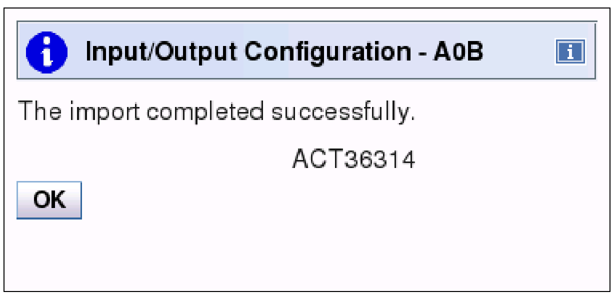

Figure 6-133 HMC - SAIOCP: source file imported

The Source Status now indicates `Imported` (Figure 6-134).

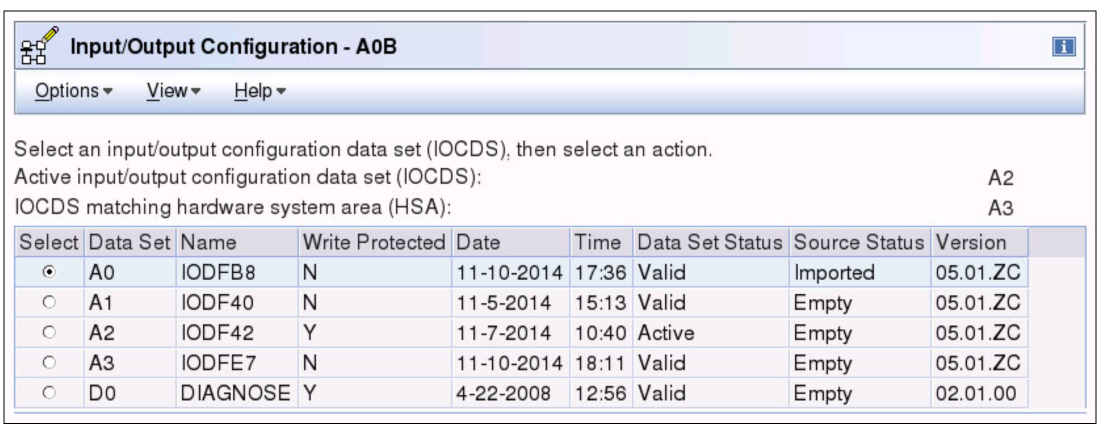

Figure 6-134 HMC - SAIOCP: IOCDS source status has changed to imported

15. Click **Options** → **Build Data Set** (Figure 6-135).

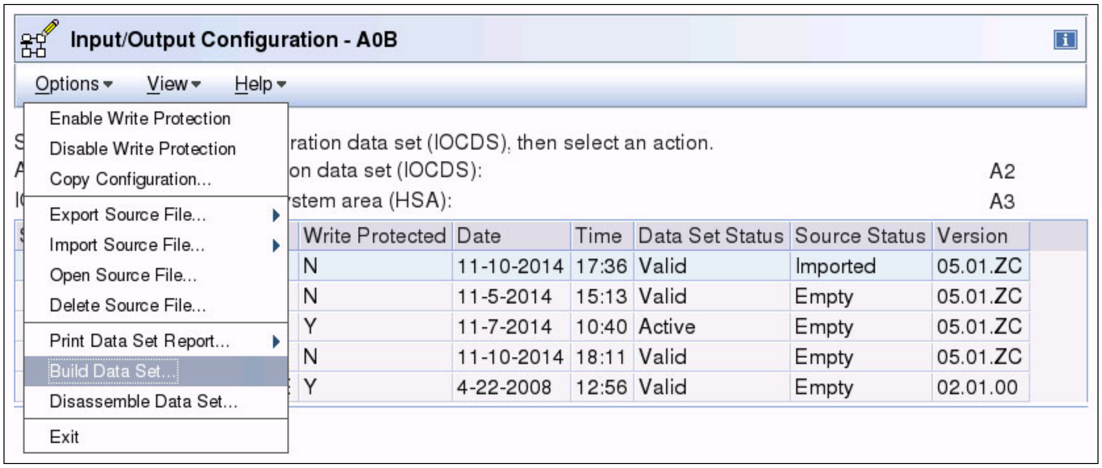

Figure 6-135 HMC - SAIOCP: building the IOCDS

16. Select the build options you want and click **OK** (Figure 6-136).

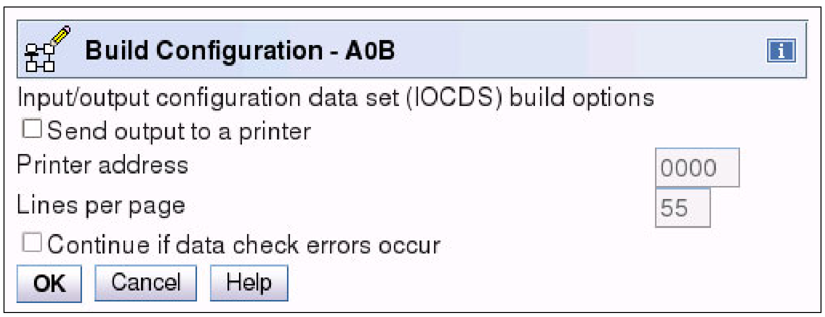

Figure 6-136 HMC - SAIOCP: build options

17. Observe the Build warning message, enter your HMC password, and click **Yes** (Figure 6-137 on page 344).

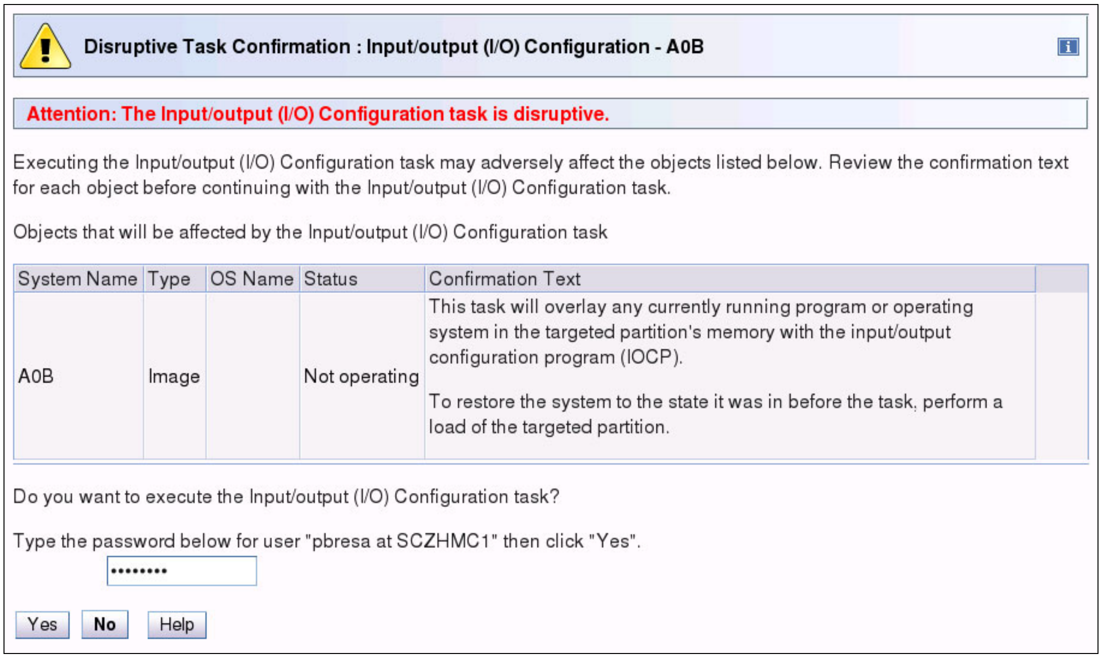

Figure 6-137 HMC - SAIOCP - build warning window

18. Status messages are displayed during the build process. After the process completes successfully, a message is displayed (Figure 6-138). Click **OK**.

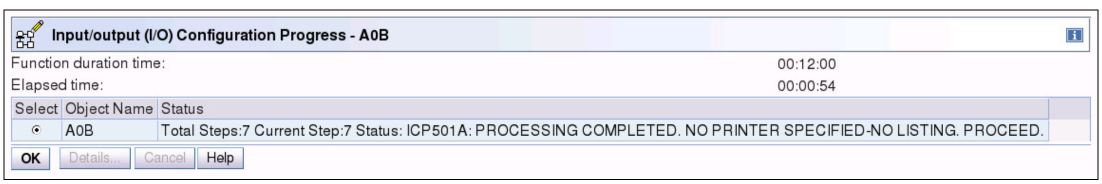

Figure 6-138 HMC - SAIOCP - build process status message

19. Observe that the Source Status now indicates `Verified` (Figure 6-139).

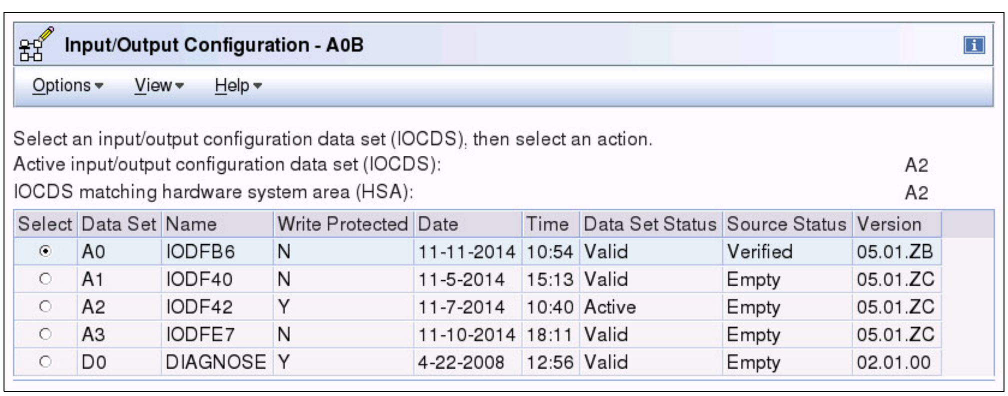

Figure 6-139 HMC - SAIOCP - IOCDS source status changed to verified

20. Click **Options** → **Exit** to end the IOCDS build process and deactivate the LPAR if it is no longer required.

This IOCDS is now ready to be selected by a Reset Profile. The 2964 can be activated (power-on reset) with the production IODF.

The USB drive can also now be removed. A drive removal message is displayed (Figure 6-140).

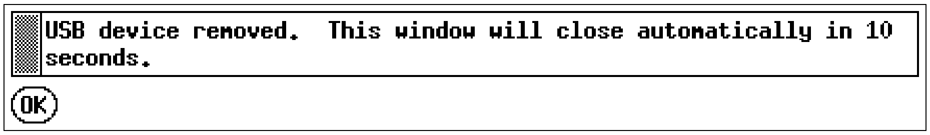

Figure 6-140 HMC - SAIOCP: USB drive removal message

6.10 HMC: Steps for profile activation

To activate your Reset Profile using the HMC, complete the steps in this section.

6.10.1 Building the Reset Profile for activation 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 2964 after it is upgraded and handed over from the IBM System Service Representative.

To build the profile, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, or use a remote web browser and select the new 2964.
2. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
3. Under Systems or Members, click the system to select it (in this example, SCZP501).
4. On the Tasks panel, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (Figure 6-141).

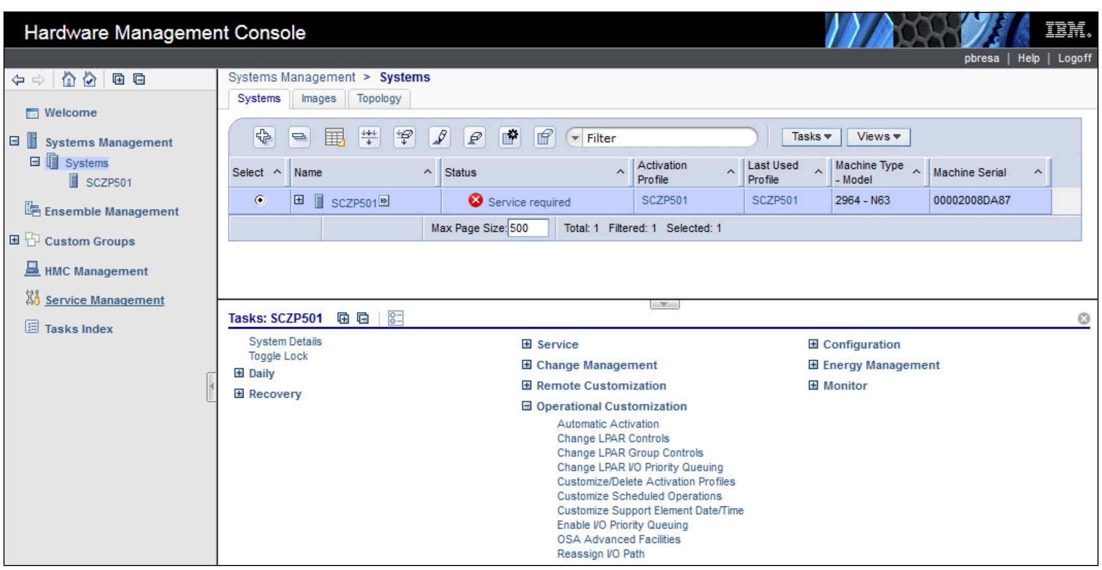

Figure 6-141 HMC - Customize Activation Profiles

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 this new TESTRESET profile and click **Customize profile**.
8. Click the IOCDS that you updated in 6.9, “HCD/HMC: Loading the 2964 processor IOCDS” on page 334. The ACTB0PDL message is displayed (Figure 6-142).

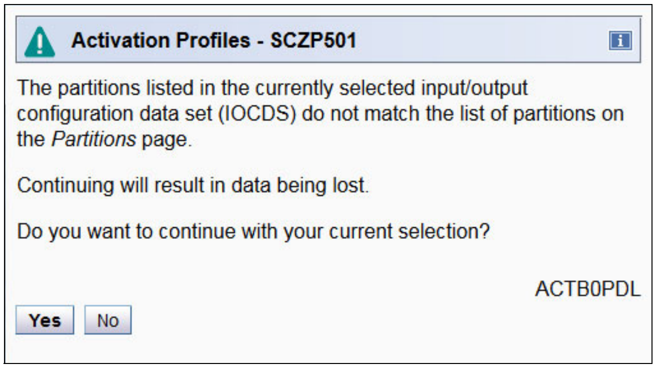

Figure 6-142 HMC - Activation Profiles: ACTB0PDL message

9. Depending on the circumstances, you can answer **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 in the IOCDS that was selected. It also allows you to create new Image profiles for those that it cannot retrieve. In the example, the last option is selected. Click **OK** (Figure 6-143).

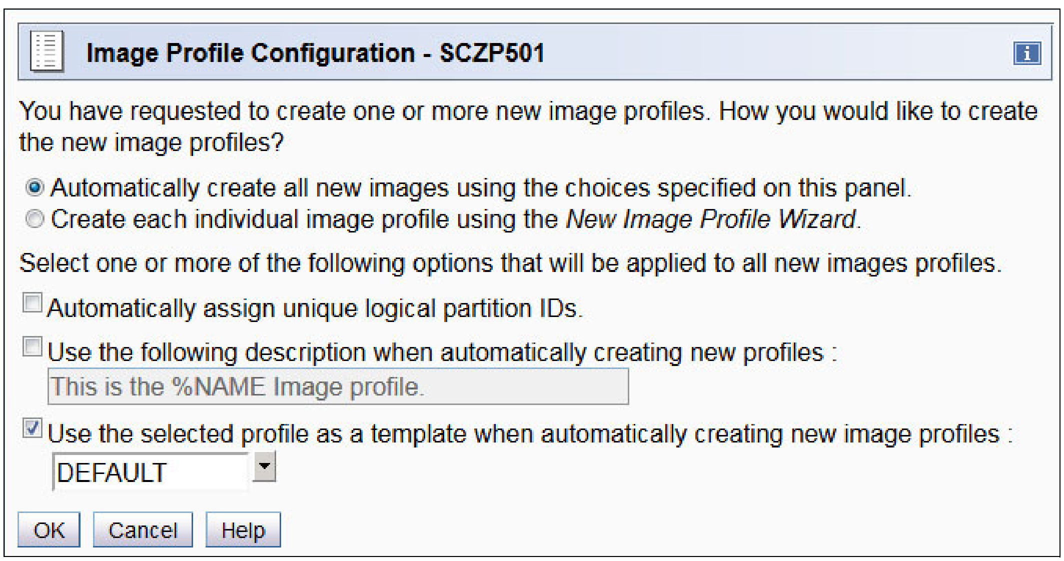

Figure 6-143 HMC - Image Profile automatic build options

11. Note the list of LPARs that were retrieved and built based on the LPARs that were defined in the selected IOCDS. Click **Save** (Figure 6-144).

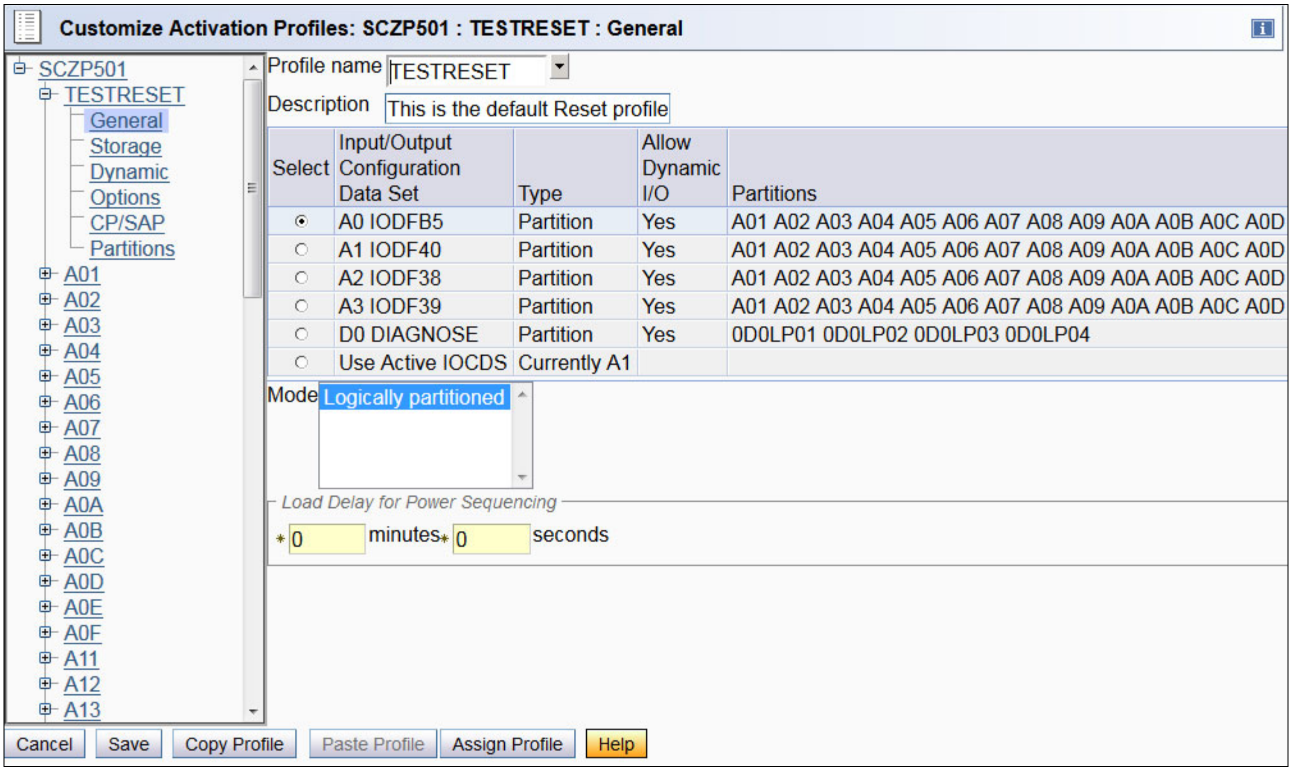

Figure 6-144 HMC - Reset and Image profile list

12. Review the items mentioned in these sections:

– 6.10.2, “Building and verifying Image Profiles” on page 347
– 6.10.5, “Server Time Protocol (STP) configuration” on page 348

6.10.2 Building and verifying Image Profiles

While still in the Reset Profile, you can review and update the Image Profile attributes.

6.10.3 Building and verifying Load Profiles

Create Load (IPL) Profiles using the DEFAULTLOAD Load profile as a template for all the logical partitions for which you are performing an IPL on this processor.

6.10.4 Building and verifying Loadxx Members in SYS#.IPLPARM

You might need more Loadxx members defined in `SYS#.IPLPARM` for this processor.

If you used the HWNAME parameter to point to the Processor ID, update this parameter to point to the new Processor ID (in this example, from SCZP401 to SCZP501). Sometimes the LPARNAME parameter is also used in the Loadxx members and might need to be reviewed or updated.

6.10.5 Server Time Protocol (STP) configuration

Review the Server Time Protocol (STP) configuration to ensure that the correct External Time Source (ETS) was configured, and that you are connected to the correct Coordinated Time Network (CTN).

For more information about setting up your STP environment, see Chapter 9, “Server Time Protocol (STP)” on page 511.

6.10.6 Running a power-on reset (POR) of the 2964

When the 2817 or 2827 processor is upgraded with a 2964, the IBM System Service Representative will perform 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 in 6.10.1, “Building the Reset Profile for activation and pointing to required IOCDS” on page 345.

The 2964 is now ready to be activated (power-on reset) by using the Production Reset Profile.

7

Installing a new IBM z13

This chapter describes how to install an IBM z13 into a new hardware environment.

Your environment will probably not contain the same elements as the configuration described here. Nevertheless, the steps in this process provide enough information for you to replicate the approach in your own environment.

This chapter includes the following topics:

7.1 Scenario overview

This section describes the new installation scenario.

7.1.1 The configuration process

The ten I/O configuration steps that are described in the “I/O configuration process” topic of *I/O Configuration Using z/OS HCD and HCM,* SG24-7804 are used for the example scenario.

Figure 7-1 shows the general process flow that is followed in this example. The numbered steps are described after the figure.

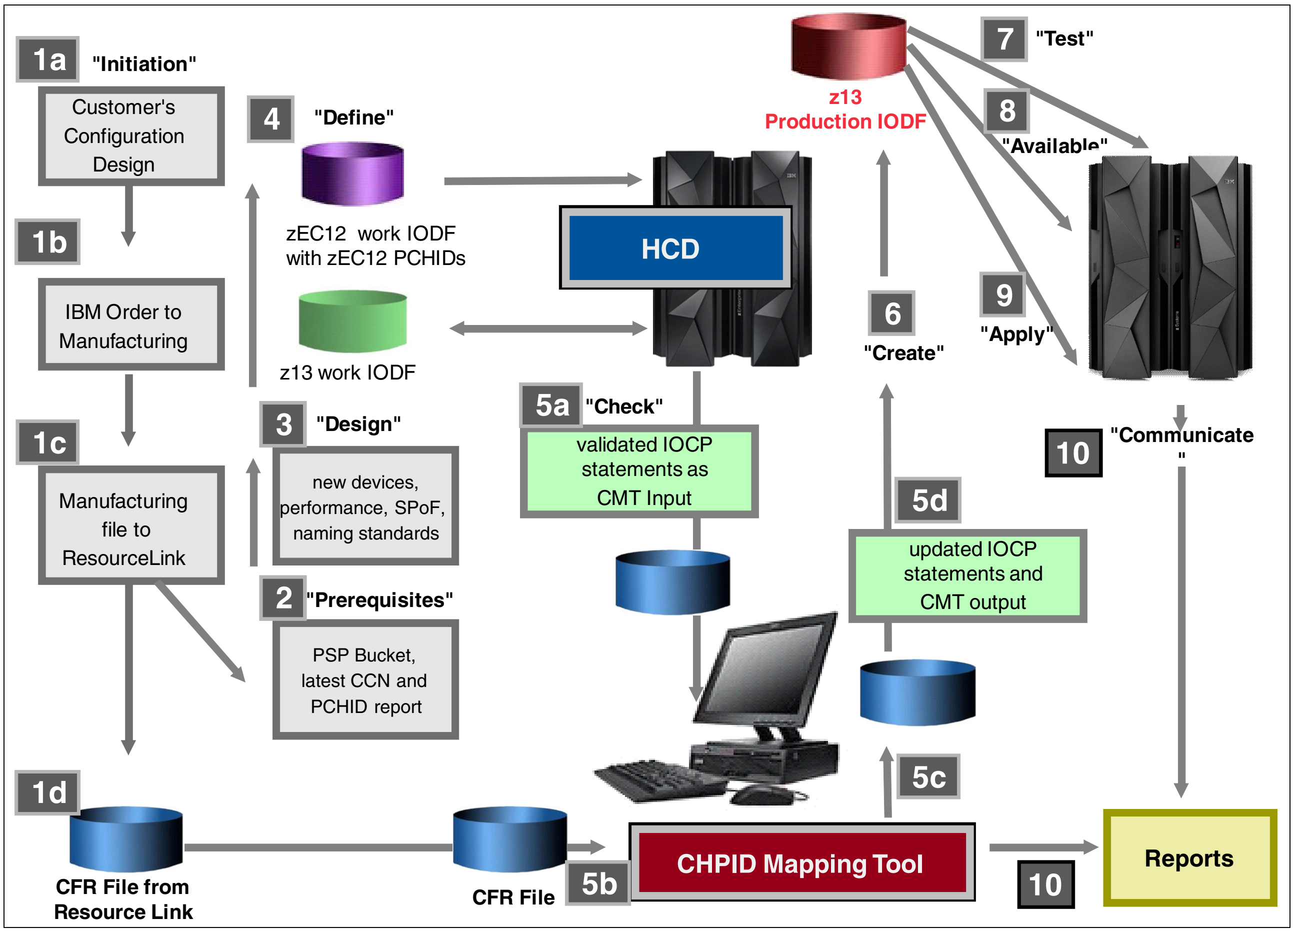

Figure 7-1 Overall configuration process flow

The steps in the figure are as follows:

1. Initiation
   a. When planning to migrate to an IBM z13, the IBM Technical Support team can help you define a configuration design that meets your needs. The configuration is then used during the ordering process.
   b. The IBM order for the configuration is created and passed to the manufacturing process.

   c. The manufacturing process creates a configuration file that is stored at the IBM Resource Link website. This configuration file describes the hardware being ordered. This data is available for download by the client installation team.
   d. A New Order report is created that shows the configuration summary of what is being ordered along with the Customer Control Number (CCN). The CNN can be used to retrieve `CFReport` (a data file that contains a listing of hardware configuration and changes for a central processor complex (CPC)) from Resource Link.
2. Prerequisites

   Ensure that you have the current PSP Bucket installed. Also, run the SMP/E report with fix category (FIXCAT) exceptions to determine whether any program temporary fixes (PTFs) must be applied. Ensure that you have the most current physical channel ID (PCHID) report and CCN from your IBM System Service Representative.
3. Design

   When you plan your configuration, consider this information:
   – Naming standards
   – FICON switch and port redundancy
   – Adequate I/O paths to your devices for performance
   – OSA channel-path identifier (CHPID) configuration for network and console communications
   – Coupling facility connections internally and to other systems.

   Because the IBM z13 server does not support attachment to the IBM Sysplex Timer, you must consider how the IBM z13 will receive its time source. A z13 cannot join a Coordinated Time Network (CTN) that includes a z10 or earlier as a member. Because the z10 was the last server that supported the IBM Sysplex Timer (9037) connectivity, the z13 cannot be configured as a member of a mixed CTN. The z13 can only join an STP-only CTN.

   When you are planning to replace a z196 or zEC12 with a new z13, plan the replacement of channels that are not supported on IBM z13. You must carefully plan how to replace those, for instance, ISC-3 to HCA3-O or ICA SR for connectivity between z13 and another z13.

   You might need to increase CF storage size when you replace z196 or zEC12 with z13. Coupling Facility Control Code (CFCC) level 20 requirements might differ from CFCC level 19 and earlier. Use the CFSizer Tool to get the new CF storage requirements.
4. Define

   The existing z196 or zEC12 I/O configuration is used as a starting point for using Hardware Configuration Definition (HCD). The z196 or zEC12 production input/output definition file (IODF) is used as input to HCD to create a work IODF that becomes the base of the new IBM z13 configuration.

   When the new IBM z13 configuration is added and the obsolete hardware is deleted, a validated version of the configuration is saved in an IBM z13 validated work IODF.
5. Check
   a. From the validated work IODF, create a file that contains the IBM z13 IOCP statements. This IOCP statements file is transferred to the workstation used for the CHPID Mapping Tool (CMT). Hardware Configuration Manager (HCM) can also be used here to transfer the IOCP deck to and from the CMT.
   b. The configuration file that is created by the IBM Manufacturing process in step 1d is downloaded from Resource Link to the CMT workstation.

      The CHPID Mapping Tool (CMT) uses the input data from the files to map logical channels to physical ones on the new IBM z13 hardware.

You might have to make decisions in response to the following situations, among others:

i. Resolving situations in which the limitations on the purchased hardware cause a single point of failure (SPoF). You might must purchase more hardware to resolve some SPoF situations.

ii. Prioritizing certain hardware items over others.

c. After the CMT processing finishes, the IOCP statements contain the physical channels to logical channels assignment that is based on the actual purchased hardware configuration.

The CHPID Mapping Tool (CMT) also creates configuration reports to be used by the IBM System Service Representative and the installation team.

The file that contains the updated IOCP statements created by the CMT, which now contains the physical channels assignment, is transferred to the host system.

d. Use HCD, the validated work IODF file created in step 5a, and the IOCP statements updated by the CMT to apply the physical channel assignments created by the CMT to the configuration data in the work IODF.

6. Create

After the physical channel data is migrated into the work IODF, an IBM z13 production IODF is created and the final IOCP statements can be generated. The installation team uses the configuration data from the IBM z13 production IODF when the final power-on reset is done, yielding a IBM z13 with an I/O configuration ready to be used.

7. Test

IODFs that are modifying existing configurations can be tested in most cases to verify that the IODF is making the intended changes.

8. Available

a. If you are upgrading an existing z196 or zEC12, you might be able to use HCD to write an IOCDS to your system in preparation for the upgrade. If you can write an IOCDS to your current system in preparation for upgrade, do so and let the IBM System Service Representative know which IOCDS to use.

> **Tip:** Using the HCD option “Write IOCDS” in preparation of an upgrade is the preferred method for writing the initial IOCDS when upgrading from a z196 or a zEC12 to an IBM z13. This scenario uses the HCD option *Write IOCDS process*.

b. If the z196 or zEC12 is not network connected to the CPC where HCD is running, or if you are not upgrading or cannot write an IOCDS in preparation for the upgrade, use HCD to produce an IOCP input file. Download this input file to a USB flash drive.

9. Apply

The new production IODF can be applied to the IBM z13 in these ways:

– Using the power-on reset process
– Using the Dynamic IODF Activate process

10. Communicate

Communicating new and changed configurations to operations and the appropriate users and departments is important.

7.1.2 Planning considerations

The following I/O features can be ordered for a new IBM z13:

- FICON Express16S LX (long wavelength - 10km) - FC 0418
- FICON Express16S SX (short wavelength) - FC 0419
- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Crypto Express5S - FC 0890
- Flash Express - FC 0403
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171
- PCIe-O SR Integrated Coupling Adapter - FC 0172

The following features, if present in a z196 or zEC12, can be carried forward when you upgrade to an IBM z13:

- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- FICON Express8 LX (long wavelength - 10 km) - FC 3325
- FICON Express8 SX (short wavelength) - FC 3326
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- OSA-Express4S GbE LX (long wavelength) - FC 0404
- OSA-Express4S GbE SX (short wavelength) - FC0405
- OSA-Express4S 10 GbE LR (long reach) - FC 0406
- OSA-Express4S 10 GbE SR (short reach) - FC 0407
- OSA-Express4S 1000BASE-T Ethernet - FC 0408
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Flash Express - FC 0402
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171

The following features are not supported on IBM z13:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet

- OSA-Express2 10 GbE Long Reach
- OSA-Express
- Crypto Express3
- Crypto Express2
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode
- ICB-2
- ICB-3
- ICB-4
- PCIXCC and PCICA
- Parallel channels (use FICON and ESCON converters)

Table 7-1 lists the channel types as described in an input/output configuration data set (IOCDS) that are supported by the z196 (2817), zEC12 (2827), and z13 (2964).

Table 7-1 Channels, links, and adapters with CHPID type and support

Channels	CHPID type	2817 support	2827 support	2964 support
ESCON channels: Connection Channel (ESCON architecture) Channel to Channel (connects to CNC) Converter Channel Path (for BL types) Converter Channel Path (for BY types)	CNC CTC CVC CBY	Up to 240 Up to 240 Up to 240 Up to 240	- No - -	- No - -
FICON native channels that attach to FICON directors or directly to FICON control units: FICON Express 4 SX and LX FICON Express 8 SX and LX FICON Express 8S SX and LX FICON Express 16S SX and LX	FC FC FC FC	Carry forward up to 288 Yes No	Carry forward Carry forward up to 320 No	No up to 64 Carry Forward up to 320
FICON channels that attach to Fibre Channel devices, switches, directors, or Fibre-Channel-to-SCSI bridges	FCP	Yes	Yes	Yes
ISC-3 peer mode channels (connects to another CFP)	CFP	Up to 48	Carry forward	No
IC peer channels (connects to another ICP)	ICP	Up to 32	Up to 32	Up to 32
HCA2-O 12x InfiniBand host channel adapters **HCA2-O 1x LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 32	Up to 32 links Up to 32 links	No No
HCA3-O 12x InfiniBand host channel adapters **HCA3-O LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 48	Up to 32 links Up to 64 links	Up to 32 links Up to 64 links
PCIe-O SR Integrated Coupling Adapter	CS5	No	No	Up to 32 links
HiperSockets (IQDIO) channels	IQD	Up to 32	Up to 32	Up to 32
OSA- Express2 GbE LX/SX	OSD and OSN	Up to 48 ports carried forward	No	No
OSA- Express2 1000BASE-T	OSE, OSD, OSC, and OSN	Yes	No	No
OSA- Express2 10 GbE LR	OSD	Yes	No	No

Channels	CHPID type	2817 support	2827 support	2964 support
OSA-Express3 GbE LX/SX	OSD and OSN	Up to 96 ports	Carry forward	No
OSA-Express3 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	OSC, OSD, OSE, and OSN: Yes OSM: No	Carry forward	No
OSA- Express3 10 GbE LR/SR	OSD and OSX	OSD: Yes OSX: No	Carry forward	No
OSA-Express4S GbE LX/SX	OSD	Up to 96 ports	Up to 96 ports	Carry forward
OSA-Express4S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Carry forward
OSA-Express4S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Carry forward
OSA-Express5S GbE LX/SX	OSD	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Up to 48 ports
10GbE RoCE Express SR	ROCE	No	Up to 32 ports	Up to 32 ports
zEDC Express Enterprise Data Compression	ZEDC	No	Max 8 features	Max 8 features

Consider the following information when you plan your configuration:

- Coupling links
- Hardware Management Console (HMC)
- Software support
- Open Systems Adapter - Integrated Console Controller
- Fibre Channel Protocol
- CPC name versus Processor ID
- Local system name

Coupling links

Only the following coupling facility CHPIDs are supported:

- CHPID type CIB: PSIFB links connecting to an HCA3-O 12x and HCA3-O LR 1x
- CHPID type ICP: Internal Coupling links
- CHPID type CS5: Integrated Coupling adapter PCIe-O SR

Considerations: Coupling links can be defined as both Coupling and STP links, or STP-only links. IBM z13 does not support the ISC-3 features.

Hardware Management Console (HMC)

The HMC can appear either as the current HMC does, or as an HMC that can run code to manage an Ensemble. The current HMC is used to manage, monitor, and operate one or more IBM z Systems servers and their associated logical partitions. An HMC that has Ensemble code running is an HMC attached to one or more zEnterprise Systems configured as Ensemble members. A particular Ensemble is managed by a pair of HMCs in primary and alternate roles.

The HMC has a global (Ensemble) management function, whereas the SE has local node management responsibility. When tasks are performed on the HMC, the commands are sent to one or more Support Elements (SEs), which then issue commands to their CPCs.

The IBM z13 requires HMC Application V2.13.0 (driver level 21) or later, and uses only Ethernet for its network connection. The HMC and the SEs do not contain a floppy disk drive. Therefore, you must use a USB flash memory drive to input and back up client configuration data.

Software support

HCD V2.1 and later with the Preventive Service Planning (PSP) bucket for 2964DEVICE and PTFs) is required to define and support some of the new features of the IBM z13.

Open Systems Adapter - Integrated Console Controller

Because support has been withdrawn for the 2074 console controllers, consider using OSA-Express4S 1000BASE-T or OSA-Express5S1000BASE-T CHPIDs defined as TYPE=OSC. With this OSA card, you can set up console function that is supported by a configuration file defined on the Support Element for that processor.

Fibre Channel Protocol

When you use CHPIDs defined as TYPE=FCP, consider N-Port ID Virtualization (NPIV).

For more information about FCP CHPIDs and the WWPN prediction tool to manage them, see “Batch Network Analyzer and Compression” on page 47.

CPC name versus Processor ID

HCD allows you to define different processors (logical) to the same CPC (physical). The Processor ID must be unique within the same IODF, but the CPC name does not. Therefore, the CPC name does not need to match the Processor ID. This advantage is useful when you have several processor/logical partition/control unit setups that share a physical CPC within the same IODF. Furthermore, the Processor ID is what is defined for the optional HWNAME parameter in the LOAD member in SYS1.IPLPARM.

The CPC name is coded in HCD option 1.3 under View Processor Definition in the CPC name field under SNA address, along with a Network name. It is the CPC name, and not the Processor ID, that is displayed on the HMC.

When you view the network information for a CPC through the HMC, the SNA address is made up of a network name and CPC name separated by a dot (for example, `USIBMSC.SCZP501`). These values are defined in the Support Element for the CPC. They must match the values that are set in the IODF so that HCD option 2.11 can find the CPC to write an IOCDS in the System z Cluster List.

Local system name

An extra system name, LSYSTEM, is used to identify the local system name of a server when you define PSIFB type=CIB or ICA type=CS5 coupling links.

This data field can be found when you change a CIB-capable processor under HCD option 1.3.

The LSYSTEM field can be set or changed to any one to eight alphanumeric characters. Also, it can begin with either an alphabetic or numeric character. All characters are uppercase.

The following rules determine whether, and where, HCD sets the LSYSTEM keyword automatically:

- If a CIB-capable processor is defined and the CPC name is set but the local system name is not set, HCD sets the local system name to the CPC name.
- If a CIB-capable processor that has not yet defined a CPC name is changed to obtain a CPC name but no local system name, HCD sets the CPC name to the local system name.
- If a non-CIB capable processor is changed to a support level that is CIB capable, and the processor has a CPC name set but no local system name, the local system name defaults to the CPC name.
- If the processor definition is changed such that the local system name is explicitly removed, HCD does not do any defaulting.
- If a processor has a local system name set (whether it has a CPC name or not), any change to the local system name must be done explicitly. There is no implicit name if the CPC name or the support level is changed.
- During Build Production IODF, verification determines whether a local system name is set for the processor if the processor has a CIB channel path defined. If this verification fails, an error message is displayed and the production IODF is not built.

Generally, set the local system name the same as the CPC name.

The following additional keywords are for the ID statement in an IOCP deck:

CPATH	Specifies the CSSID and CHPID on the connecting system.
CSYSTEM	Connecting System name (LSYSTEM in processor definition).
AID	Host Card Adapter ID (found in PCHID report).
PORT	Host Card Adapter port.

7.1.3 Adding a new 2964 to a new hardware environment

This scenario shows the configuration steps for installing a new 2964 system into a new environment. The 2964 is the first system on the floor and is being connected to new switches and control unit interfaces. The configuration has these key considerations:

- HCD requires a new Processor ID for the 2964.
- HCD requires a new CPC name for the 2964.
- The 2964 processor connects to new switch ports and new control unit interfaces.
- The control unit interfaces connect to new switch ports.
- The starting IODF can be an existing production IODF from another data center, or can be a new work IODF.
- The target IODF is a new 2964 work IODF.

This example shows a new 2964-N63 with a Processor ID of SCZP501 and with six CSSs (CSS ID=0, CSS ID=1, CSS ID=2, CSS ID =3, CSS ID=4, and CSS ID=5). The CPC name of SCZP501 and serial number of 02-8DA87 are used for the 2964.

The following CHPID types are defined:

- OSD, OSC, OSM, and OSX
- FC and FCP
- CIB and ICP
- IQD

The following hardware/CHPID types are *not* supported for the 2964:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- Crypto Express3
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode

Table 7-2 summarizes the tool requirements. The process is documented later in this chapter.

Table 7-2 New installation of a 2964 in a new installation

New 2964	**New 2964 to connect to new switch ports and control units that do not currently exist (new installation)**
Processor ID	Requires a new Processor ID
CPC name	Requires a new CPC name
Channel to switch port connections	New ports
Control Unit to switch port connections	New ports
Starting IODF	Current active production IODF
Target IODF	Create a work IODF
HCD action	Add a processor (see step 2 in 7.2, “HCD: Creating a 2964 IODF” on page 359)
CMT Program (needed or not)	Optional, but generally a good idea
CFReport File (CCN)	Required
IOCP (import from validated work IODF)	Yes
CMT Actions (PCHID reset)	Yes
CMT IOCP Output	Yes
CMT Reports	Yes (CHPID Report and CHPID to CU Report)

Tip: As mentioned, you can have different Processor ID and CPC names. However, if you do not need to support multiple processor images, be sure that the Processor ID and CPC names match.

7.2 HCD: Creating a 2964 IODF

The following steps explain how to define a new 2964 configuration by using HCD:

1. Create a 2964 work IODF, either new or from an existing production IODF.
2. Add the processor.
3. Change the required partition names and usage from reserved.
4. Add CHPIDs.
5. Add FICON switches.
6. Add operating system configurations.
7. Create the eligible device table (EDT).
8. Add esoteric names in the EDT.
9. Connect FICON CHPIDs to new switches, spreading the connections over as many switches and physical port cards as possible, where appropriate.
10. Create control units unique to this processor.
11. Create devices appropriate to the control units.
12. Define devices to the operating system configurations and any esoterics, where appropriate.
13. Connect control units to CHPIDs or switch ports, and then to CHPIDs.
14. Define all required coupling connections to other processors and any Internal coupling connections required.
15. Create FICON channel-to-channel connections (FCTCs).
16. Create the OSA configuration (OSC, OSD, OSE, OSN, OSX, and OSM).
17. Define nucleus initialization program (NIP) consoles.
18. Build a validated Work IODF.
19. Create an IOCP statements file and transfer it to your CHPID Mapping Tool workstation. This step can also be performed using HCM.
20. Import CFReport and IOCP statements into the CMT.
21. Perform hardware resolution and PCHID/CHPID availability.
22. Create configuration reports for yourself and the IBM System Service Representative.
23. Import IOCP statements that have been updated with PCHIDs back into a validated work IODF.
24. Build the Production IODF.
25. Remotely write the IOCP to an IOCDS on the 2964 Support Element. If this action is not possible, copy the IOCP statements to a USB memory flash drive and run the Stand-Alone Input/Output Configuration Program to load the IOCP statements to an IOCDS on the 2964 Support Element.
26. Build Reset, Image, and Load profiles if required.
27. Perform a power-on reset (activate) of 2964.

The following sections describe some of these steps in more detail. For more information about defining processors, partitions, switches, control units, and devices in an IODF by using HCD, see *z/OS HCD Users Guide*, SC33-7988 and *IOCP Users Guide*, SB10-7037.

7.2.1 Creating a work IODF

Create and select a new work IODF where you will add the new 2964 processor definition (for example, `SYS6.IODFB7.WORK`).

7.2.2 Adding a 2964 processor

To add a 2964 processor, complete the following steps:

1. From the HCD main menu, select option **1.3. Processor List**.
2. In the Processor List (Figure 7-2), press PF11, or enter **`add`** on the command line to add a processor and press Enter.

```
                              Processor List
Command ===> add_______________________________________ Scroll ===> CSR

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
****************************** Bottom of data *******************************
```

Figure 7-2 HCD - Processor List: adding a processor

The Add Processor panel opens (Figure 7-3).

```
------------------------------ 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  . . . . . . . ________
```

Figure 7-3 HCD - Add Processor: blank data fields

3. Enter the appropriate information in the Add Processor panel. For the example, specify the following values, as shown in Figure 7-4:

Processor ID	`SCZP501`
Processor type	`2964`
Processor model	`N63`
Serial number	`08DA872964`
Network name	`USIBMSC`
CPC name	`SCZP501`

```
------------------------------ Add Processor ------------------------


Specify or revise the following values.

Processor ID . . . . . . . . . . SCZP501_
Processor type . . . . . . . . . 2964____  +
Processor model  . . . . . . . . N63_____  +
Configuration mode . . . . . . . LPAR      +
Number of channel subsystems . . _         +

Serial number  . . . . . . . . . 08DA872964
Description  . . . . . . . . . . ________________________________

Specify SNA address only if part of a System z cluster:

Network name . . . . . . . . . . USIBMSC_  +
CPC name . . . . . . . . . . . . SCZP501_  +

Local system name  . . . . . . . ________
```

Figure 7-4 HCD - Add Processor: updating data fields

4. Press Enter. The Processor List panel opens (Figure 7-5 on page 362), showing the new 2964 processor, named SCZP501.

   The message displayed at the bottom of the panel indicates that the processor definition is extended to its maximum configuration. This is because part of the main storage is allocated as a fixed-sized Hardware System Area, which is not addressable by application programs. In HCD, when you define as new or redefine a processor as a 2964, HCD automatically defines the maximum configuration of six CSSs and 84 logical partitions.

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

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


 / Proc. ID Type +   Model +  Mode+ Serial-# + Description
 _ SCZP501  2964     N63      LPAR  08DA872964 ________________________________
 ****************************** Bottom of data *******************************




+-------------------------------------------------------------------+
| Definition of processor SCZP501 has been extended to its maximum |
| configuration.                                                    |
+-------------------------------------------------------------------+
```

Figure 7-5 HCD - Processor List: new processor definition

5. Enter `s` next to the SCZP501 and press Enter. The Channel Subsystem List is displayed. Here you can see six channel subsystems (CSS0-CSS5) that are defined with the default MAXDEV values for SS0 of 65280 set by HCD and 65535 set for SS1, SS2 and SS3 (Figure 7-6 on page 363).

Tip: The z13 and the zEC12 support IPL from subchannel set 1 (SS1), subchannel set 2 (SS2), or subchannel set 3 (SS3) in addition to subchannel set 0. 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, defined using the same device number and a new device type in an alternate subchannel set, to be used for IPL, I/O definition file (IODF), and stand-alone dump volumes when needed.

IPL from an alternate 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.

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

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

Processor ID . . . : SCZP501

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

Figure 7-6 HCD - Channel Subsystem List: four subchannel sets

6. Enter **p** next to one more CSS and observe that HCD also defines the maximum number of logical partitions, 15 per CSS (and 10 in CSS5) to make a total of 85, as Reserved (*), as shown in Figure 7-7.

```
---------------------------- Partition List ----------------------------
  Goto  Backup  Query  Help
-----------------------------------------------------------------------
                                                              Row 1 of 15
Command ===> ________________________________________ Scroll ===> CSR

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

Processor ID  . . . . : SCZP501
Configuration mode  . : LPAR
Channel Subsystem ID  : 0

/ Partition Name   Number Usage + Description
_ *                1      CF/OS   ________________________________
_ *                2      CF/OS   ________________________________
_ *                3      CF/OS   ________________________________
_ *                4      CF/OS   ________________________________
_ *                5      CF/OS   ________________________________
_ *                6      CF/OS   ________________________________
_ *                7      CF/OS   ________________________________
```

Figure 7-7 HCD - Partition List

Scroll to see the remaining reserved logical partitions that are defined for this CSS.

7. Define the resources to the new 2964 processor:
   a. Change Reserved partitions for each CSS to the required partition name and usage. Not all partitions must be changed now. They can be modified later by using the Dynamic IODF Activate process.
   b. Add CHPIDs to each CSS, with no PCHIDs assigned.
   c. Define a Partition Access list for these CHPIDs.
   d. Define a Partition Candidate list for these CHPIDs.

More information: For more information about defining processors, partitions, switches, control units, and devices, see the *z/OS HCD Users Guide,* SC33-7988, and the *IOCP Users Guide,* SB10-7037.

8. Add FICON switches.
9. Add operating system configurations:
   a. Create an EDT for each of these operating system configurations.
   b. Add esoteric names in each of these EDTs.
10. Connect FICON CHPIDs to the new switches, spreading connections over as many switches and physical port cards as possible, where appropriate.
11. Create control units:
    a. Create devices appropriate to the control units.
    b. Define devices to operating system configurations and any esoterics, where appropriate.
    c. Connect control units direct to FICON CHPIDs or to switch ports, and then CHPIDs.
12. Define all required internal coupling links and peer coupling links to any other new processors you might be adding to the IODF.
13. Create FICON CTCs.
14. Create OSA control units (OSC, OSD, OSE, OSN, OSM, and OSX).
15. Define NIP consoles.

7.2.3 Overdefining channel paths on an XMP processor

Sometimes, you must define a channel path that is not physically installed on the processor. This definition is useful if you plan to add more channel cards to the processor in the future and want to have the definitions in the IODF before the hardware is installed.

HCD allows you to overdefine CHPIDs by using an asterisk (*) for the PCHID value. An overdefined CHPID must adhere to all validation rules, but it is not taken into account by an IOCDS download. Also, it is not included in the IOCP statements, in a CONFIGxx member, or during dynamic activation.

If a control unit contains only CHPIDs with a PCHID value of an asterisk (*), the whole control unit (including any attached devices) is omitted from the configuration to be activated.

When you install the channel path later, you must edit the CHPID and replace the asterisk (*) with its valid PCHID.

Remember: This is not the case for CIB type CHPIDs, where these CHPIDs have connections to other CIB type CHPIDs. Therefore, HCD allows you to define CIB type CHPIDs as overdefined only if they are unconnected.

Overdefining is also supported for CS5 type CHPID definitions.

The 2964 production IODF can then be activated dynamically and the PCHID, CHPID, and control unit definitions become available to the operating system.

Figure 7-8 shows the CHPID/PCHID definitions *before* they are defined as overdefined. Press PF20 (right) in the Channel Path List.

```
                              Channel Path List     Row 1 of 107 More: <
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    1A3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    553   FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    153   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 7-8 HCD - Channel Path List: reserving CHPIDs

Figure 7-9 shows the CHPID/PCHID definitions *after* they are defined as overdefined.

```
                               Channel Path List    Row 22 of 109 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    *     FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    *     FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    *     FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 7-9 HCD - Channel Path List: overdefined CHPIDs

7.3 HCD: Validating the 2964 work IODF

To validate the 2964 work IODF by using HCD, complete the following steps:

1. Select HCD option **2.12. Build validated work I/O definition file**.
2. Review the message list and correct any errors.
3. Press PF3 to continue. You receive the following message:

   `Requested action successfully processed`

4. Select HCD option **6.4. View I/O Definition File Information**. The IODF type is now indicated as `Work - Validated` (Figure 7-10).

```
--------------- View I/O Definition File Information -----------


IODF name  . . . . . . : 'SYS6.IODFB7.WORK'
IODF type  . . . . . . : Work - Validated
IODF version . . . . . : 5

Creation date  . . . . : 2014-11-10
Last update  . . . . . : 2014-11-10  13:04

Volume serial number . : IODFPK
Allocated space  . . . : 4000    (Number of 4K blocks)
Used space . . . . . . : 1982    (Number of 4K blocks)
   thereof utilized (%)  67
Activity logging . . . : No
Multi-user access  . . : No
Backup IODF name . . . :

Description  . . . . . :
```

Figure 7-10 HCD - View I/O Definition File Information: validated work IODF

7.4 Creating the IOCP file for the CHPID Mapping Tool

To create the IOCP file for the CHPID Mapping Tool, complete the following steps:

Tip: You might prefer to use HCM to create the IOCP statements file and transfer the file to your workstation. You can then start the CHPID Mapping Tool, create an updated IOCP statements file, and transfer the file back to the host.

1. Select HCD option **3 Build IOCP input data set** and press Enter (Figure 7-11).

```
 ----- 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
```

Figure 7-11 HCD - Activate or Process Configuration Data: build IOCP for SCZP501

2. HCD displays the list of available processors (Figure 7-12). Select the SCZP501 processor by typing a forward slash (/) next to it and pressing Enter.

```
------------------------- Available Processors --------------------------
                                                               Row 1 of 1
Command ===> ___________________________________________________________

Select one.

  Processor ID  Type     Model    Mode  Description
/ SCZP501       2964     N63      LPAR  Sphinx
*************************** Bottom of data *****************************
```

Figure 7-12 HCD - Available Processors: selecting a processor for IOCP file

3. HCD displays a panel where you enter information about the IOCP input data set to be created (Figure 7-13).

   Complete the following fields:

   – Title1: `IODFB7`
   – IOCP input data set: `'SYS6.IODFB7.IOCPIN.SCZP501'`
   – Input to Stand-alone IOCP: `Yes`
   – Job statement information: Complete this information for your installation.

4. Press Enter. HCD submits a batch job to create the data set (Figure 7-13).

```
------------------------ Build IOCP Input Data Set ----------------------


Specify or revise the following values.

IODF name . . . . . . . . . : 'SYS6.IODFB7.WORK'
Processor ID  . . . . . . . : SCZP501
Title1 . IODFB7_________________________________________________________
Title2 : SYS6.IODFB7.WORK - 2014-11-10 13:04

IOCP input data set
'SYS6.IODFB7.IOCPIN.SCZP501'____________________________
Input to Stand-alone IOCP?  Yes  (Yes or No)

Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 7-13 HCD - Build IOCP Input Data Set: data fields to be updated

5. In Time Sharing Option (TSO), verify that the data set that you created exists and contains IOCP statements (Figure 7-14). This data set is used as input into the CHPID Mapping Tool.

```
ID    MSG1='IODFB7',                                              *
      MSG2='SYS6.IODFB7.WORK - 2014-11-10 13:04',                 *
      SYSTEM=(2964,1),LSYSTEM=SCZP501,                            *
      TOK=('SCZP501',00800001DA872964130426110114314F00000000,*
      00000000,'14-11-10','13:04:26','........','........')
RESOURCE PARTITION=((CSS(0),(A0A,A),(A0B,B),(A0C,C),(A0D,D),(A*
      0E,E),(A0F,F),(A01,1),(A02,2),(A03,3),(A04,4),(A05,5),(A*
      06,6),(A07,7),(A08,8),(A09,9)),(CSS(1),(*,1),(*,2),(*,3)*
      ,(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*
      *,D),(*,E),(*,F)),(CSS(2),(*,1),(*,2),(*,3),(*,4),(*,5),*
      (*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(**
      ,F)),(CSS(3),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*
      *,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(4),*
      (*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(**
      ,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(5),(*,1),(*,2),(*
      *,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,*
      C),(*,D),(*,E),(*,F)))
FUNCTION FID=000,VF=1,PCHID=140,PNETID=ITSOPNET1,                 *
      PART=((A01),(=)),TYPE=ROCE
FUNCTION FID=010,VF=1,PCHID=208,PNETID=ITSOPNET1,                 *
      PART=((A02),(=)),TYPE=ROCE
FUNCTION FID=011,VF=2,PCHID=208,PNETID=ITSOPNET1,                 *
      PART=((A03),(=)),TYPE=ROCE
FUNCTION FID=020,VF=1,PCHID=1BC,PART=((A01),(=)),TYPE=ZEDC
```

Figure 7-14 HCD - IOCP input data set: contents (truncated)

Part of the `TOK` statement is now replaced with dots (Example 7-1).

Example 7-1 HCD - IOCP file: TOK statement

```
TOK=('SCZP501',00800001DA872964130426110114314F00000000,*
00000000,'14-11-10','13:04:26','........','........')
```

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, which is something that can be done only for processors that contain PCHID definitions.

Important: When an IOCP statement file is exported from a validated work IODF 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 file from TSO to your workstation. Use a workstation file transfer facility, such as the one in IBM Personal Communications Workstation Program, or any equivalent 3270 emulation program. Be sure to use TEXT as the transfer type. In the example, the file is named `SCZP501in.iocp`.

7.5 CMT: Assigning PCHIDs to CHPIDs

The following steps use the output from the previous set of HCD steps (IOCP) and the 2964 order process (CFReport). Use the CHPID Mapping Tool to assign PCHIDs to each of the CHPIDs for the 2964.

For this process, the CMT must be downloaded. For more information about downloading and installing the CMT, see 3.1.5, “CHPID Mapping Tool (CMT)” on page 44. If CMT is already installed, verify that the latest updates are installed. For more information, see the *CHPID Mapping Tool User’s Guide*, GC28-6947-00.

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 PCHID conflict.
4. Process the hardware resolution.
5. Resolve manually the CIB 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 and transfer it back to the host z/OS image.

Requirement: When you upgrade from a 2817 or a 2827 to a 2964, you must use the CHPID Mapping Tool level that supports the 2964.

7.5.1 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. Right-click the Projects pane and select **New** → **Standard CMT Project** (Figure 7-15).

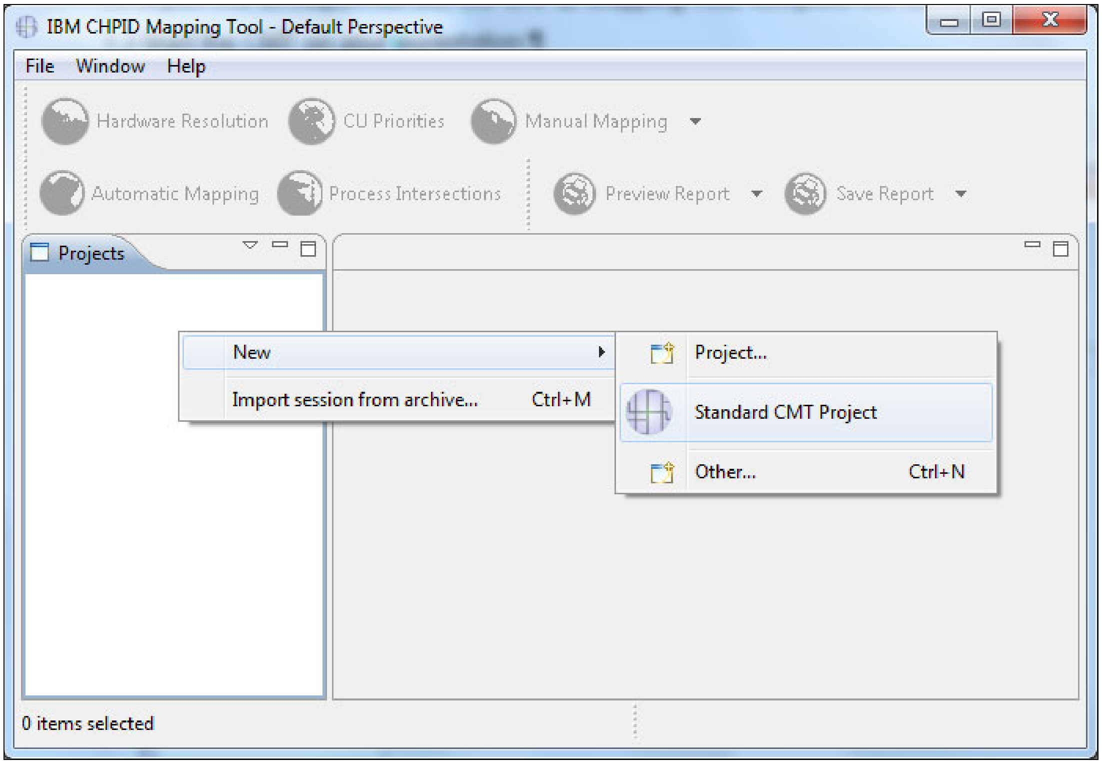

Figure 7-15 CMT - Creating a new CHPID Mapping Tool Project

3. The New CHPID Mapping Tool Project window opens (Figure 7-16). Use this window to specify a project name and then click **Next**.

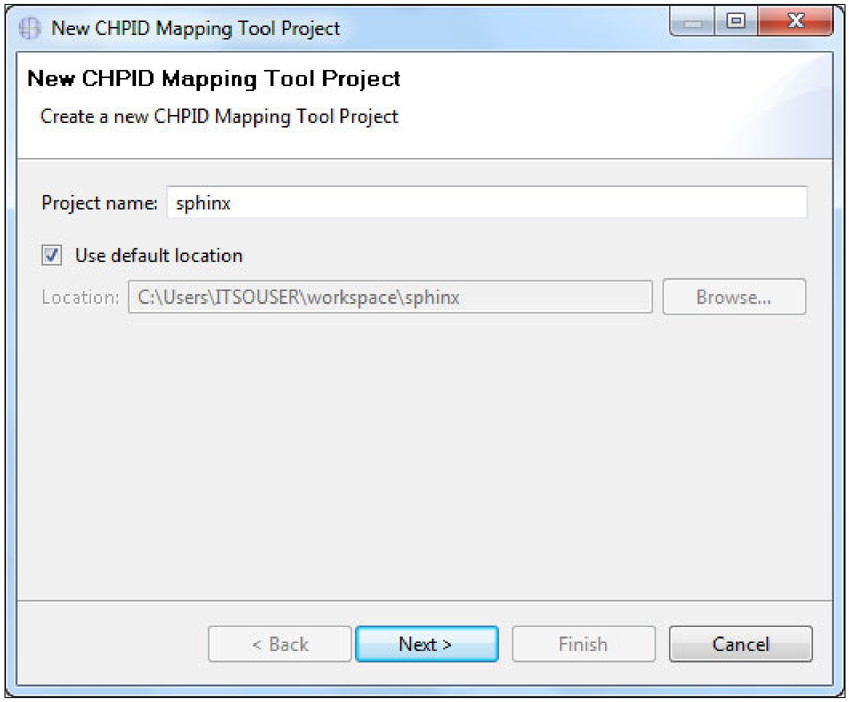

Figure 7-16 CMT - New CHPID Mapping Tool Project name

4. Import the `CFReport` file into the CHPID Mapping Tool by specifying the name in the CFReport file field, and then click **Finish** (Figure 7-17).

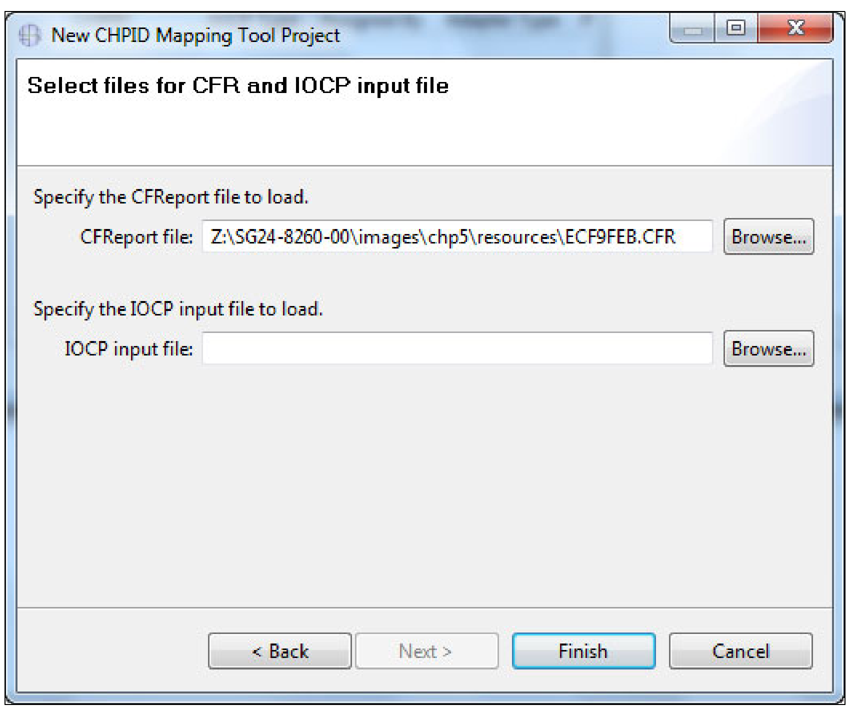

Figure 7-17 CMT - Specifying the CFReport file

If you click **Finish** but did not select an IOCP file, you receive the message shown in Figure 7-18.

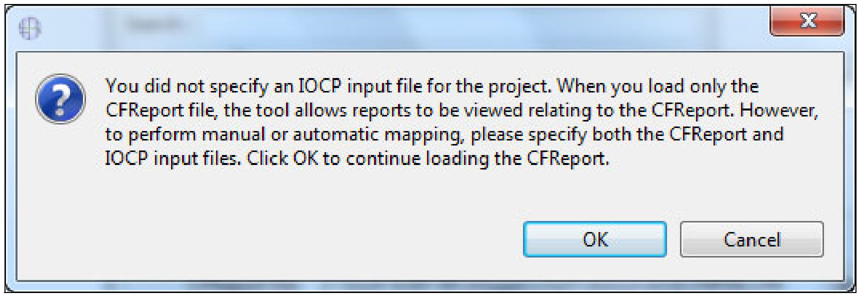

Figure 7-18 CMT - Warning message for not specifying an IOCP file

A window shows the progress of reading the CFReport file (Figure 7-19).

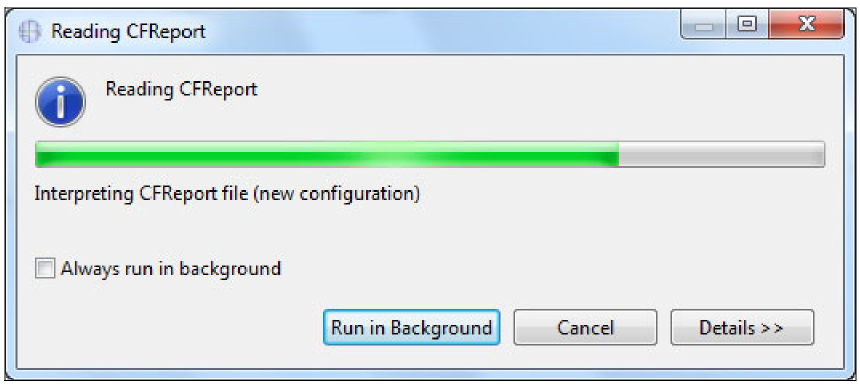

Figure 7-19 CMT - Reading the CFReport file

The information from the CFReport file is displayed in the Hardware pane (Figure 7-20).

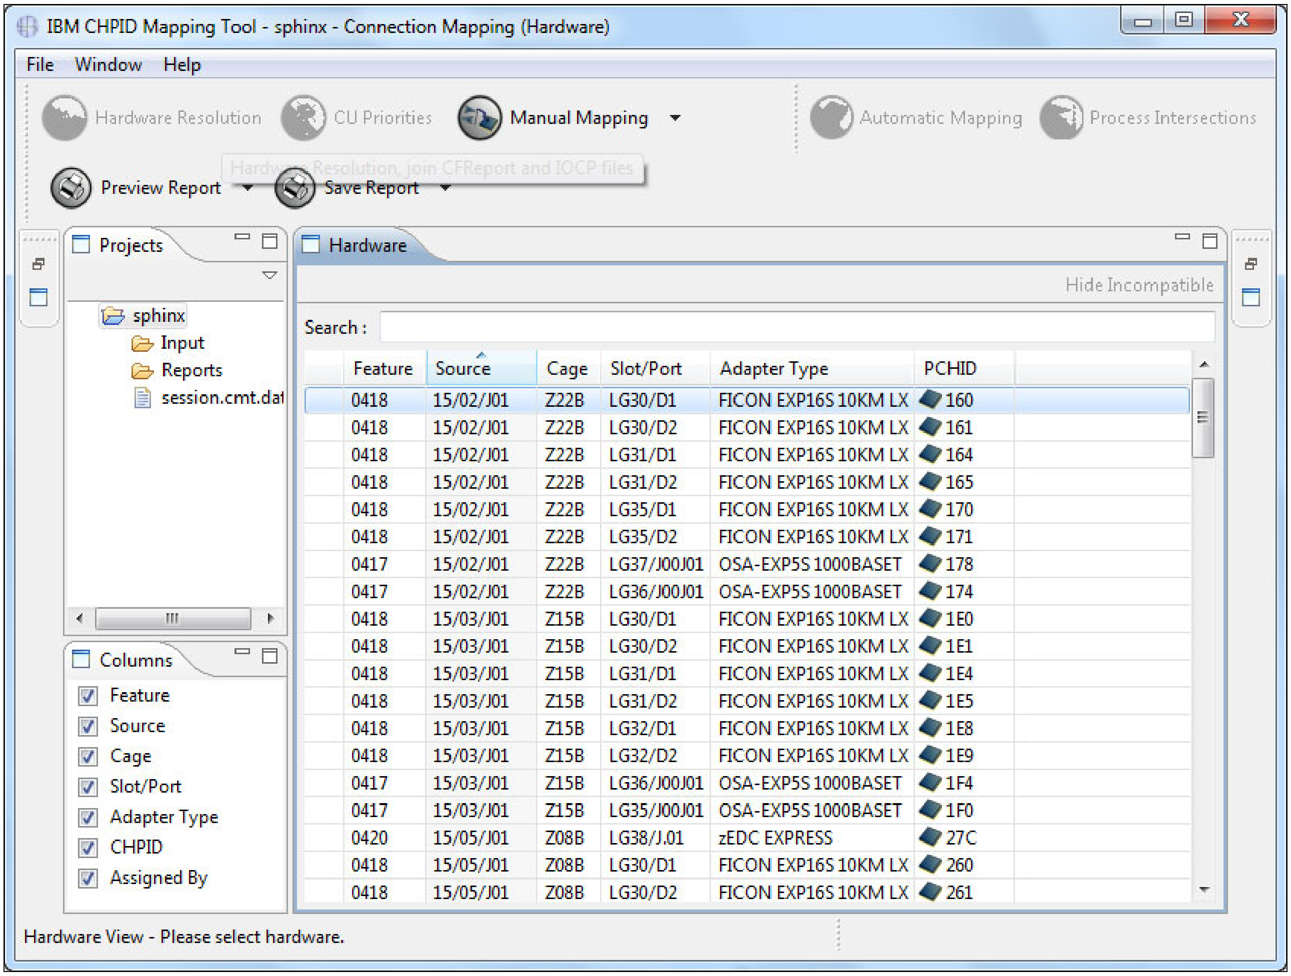

Figure 7-20 CMT - Imported CFReport file

7.5.2 Importing the 2964 IOCP file into the CHPID Mapping Tool

To import the 2964 IOCP file into the CHPID Mapping Tool, complete the following steps:

1. Right-click in the Projects pane and select **Import IOCP input file** (Figure 7-21).

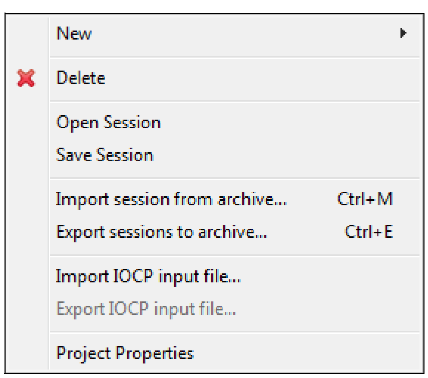

Figure 7-21 CMT - Importing the IOCP file

2. Select the IOCP file on your workstation to import into the CHPID Mapping Tool and click **Finish** (Figure 7-22).

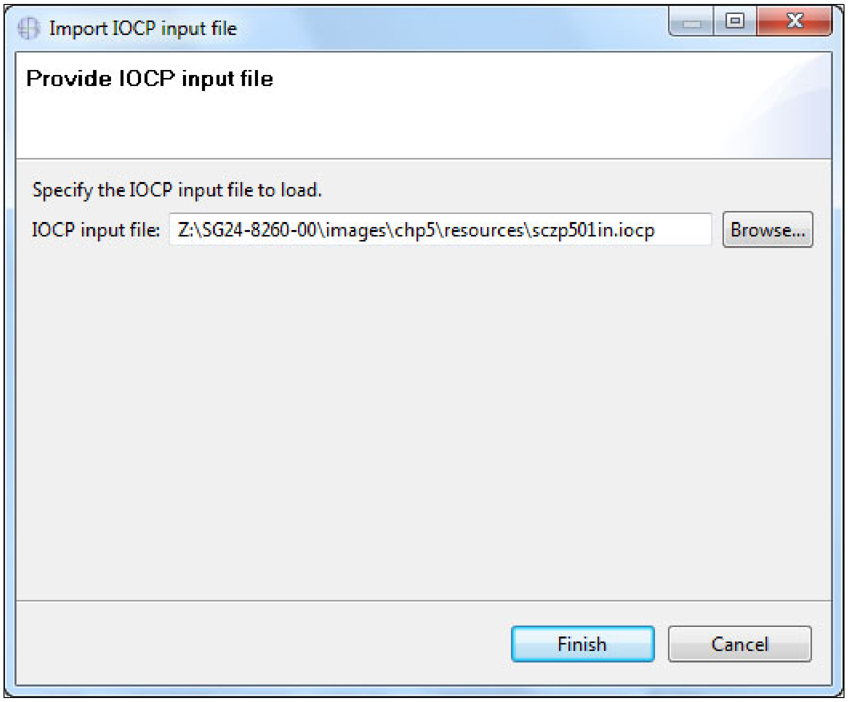

Figure 7-22 CMT - Specifying the IOCP file for import

3. In the Projects pane, under the Input tab, expand the IOCP tab, right-click the IOCP file, and select **Read Selected IOCP** (Figure 7-23).

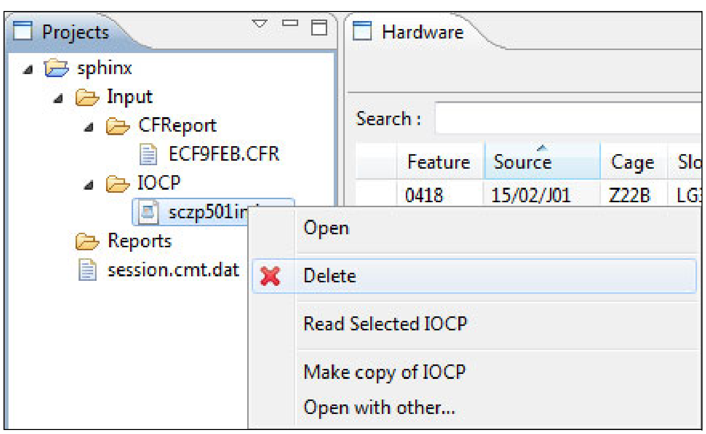

Figure 7-23 CMT - Reading the selected IOCP

A window shows the progress information (Figure 7-24).

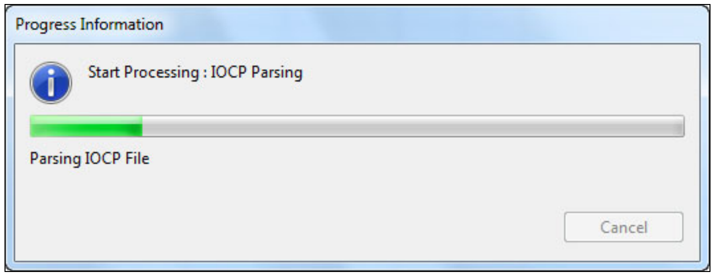

Figure 7-24   CMT - Processing the IOCP file

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 (Figure 7-25) includes three tabbed panes: Projects, Hardware Resolution, and Adapter Type Summary. Hardware Resolution is the middle pane and the Adapter Type Summary is on the right.

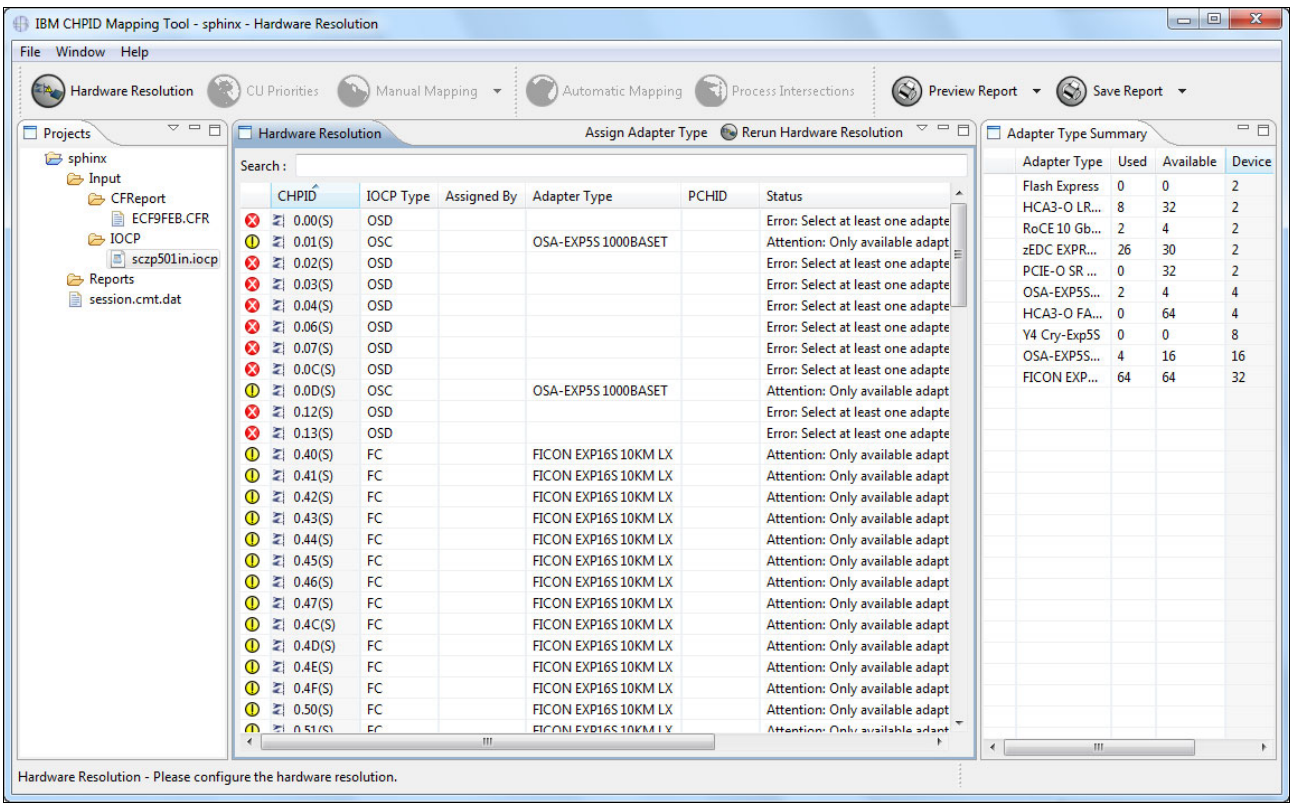

Figure 7-25   CMT - Hardware Resolution after Imported CFReport file

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 shows the following output:

- Hardware Resolution: This window lists all CHPIDs that were found; the Status column shows the CHPID information to be investigated. In the example, investigate the status; Table 7-3 lists status messages and possible resolutions.

Table 7-3 Status messages and possible resolutions

Status	Explanation	Resolution (if required)
No hardware found	AID values or PCHID values are present that are not found in the hardware. This situation could occur when you are replacing hardware for an MES and the IOCP file contains a PCHID value for the old hardware. (The IOCP file contains a PCHID value for the hardware being removed.)	If you have any CHPIDs of IOCP type CIB, 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. Do the following steps for CIB CHPIDs: 1. Remove the AID values. 2. Do one of the following tasks: - Inside the CHPID Mapping Tool, perform manual mapping to associate these CHPIDs with AIDs. - Assign the AID values outside the tool, for example, using Hardware Configuration Definition (HCD). 3. Replace the IOCP file.
Select at least one adapter type.	A adapter type is not assigned to the current row.	Assign a adapter type to IOCP type.
Adapter_type is not compatible with *IOCP_type*.	Adapter type assigned for the CHPID is not compatible with the IOCP type specified by the IOCP file.	See Performing hardware resolution for a type mismatch
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 need to change the IOCP file or obtain additional hardware
PCHID_1 moved to new channel ID: PCHID_2 **Example**: 520 moved to 1E2	You are replacing hardware for an MES, and the IOCP file contains a PCHID value for the old hardware, which is being removed. This PCHID value has moved from an old machine to the PCHID value for the new hardware. PCHID_1 is the first PCHID value (for example, 520) and PCHID_2 is the second PCHID 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 if you prefer a different assignment.

- Manual mapping CIB CHPIDs: Availability Mapping cannot be used until all CIB CHPIDs are resolved. You can use manual mapping to resolve any CIB 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 PCHID in the Manual window. If this option is not selected (it has no check mark), then availability PCHIDs 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 might require recabling after availability assignments.

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

    Regardless of whether the channels are moving or not, the CHPID Mapping Tool either assigns PCHIDs 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 PCHID where the affected CHPIDs are assigned) before you use the CHPIDs.

    By default, this option is not selected.

  If none of these options is selected, availability works only on CHPIDs that do not have PCHIDs assigned.

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

  > **Attention:** If you run **Reset CHPIDs assigned by IOCP**, it will reset any previously mapped CHPID assignments and 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 prior CHPID-to-PCHID relationship does not cause unacceptable availability.

7.5.3 Resolving CHPIDs with PCHID conflict

The CMT displays the CHPIDs with PCHID conflicts (Figure 7-26).

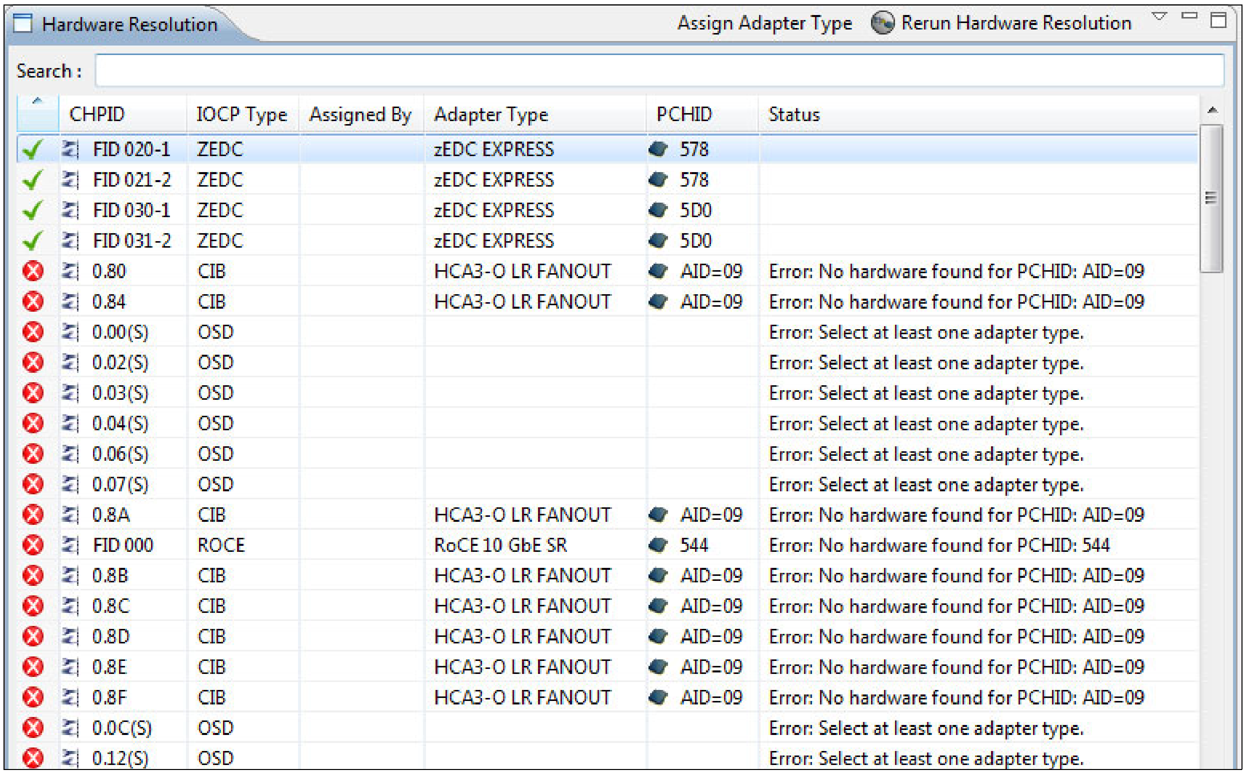
Hardware Resolution
Assign Adapter Type
Rerun Hardware Resolution
Search :

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 020-1	ZEDC		zEDC EXPRESS	578	
FID 021-2	ZEDC		zEDC EXPRESS	578	
FID 030-1	ZEDC		zEDC EXPRESS	5D0	
FID 031-2	ZEDC		zEDC EXPRESS	5D0	
0.80	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.84	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.00(S)	OSD				Error: Select at least one adapter type.
0.02(S)	OSD				Error: Select at least one adapter type.
0.03(S)	OSD				Error: Select at least one adapter type.
0.04(S)	OSD				Error: Select at least one adapter type.
0.06(S)	OSD				Error: Select at least one adapter type.
0.07(S)	OSD				Error: Select at least one adapter type.
0.8A	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
FID 000	ROCE		RoCE 10 GbE SR	544	Error: No hardware found for PCHID: 544
0.8B	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8C	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8D	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8E	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8F	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.0C(S)	OSD				Error: Select at least one adapter type.
0.12(S)	OSD				Error: Select at least one adapter type.

Figure 7-26 CMT - CHPIDs with PCHID conflicts

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

- An X in a red circle: This indicates an error.
- A green check mark: This indicates that the tool successfully resolved the specified Channel Type.

The example has these reasons to resolve hardware resolution issues:

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

7.5.4 Hardware Resolution

In the example, the CHPID Mapping Tool displays an X in the first column of the Hardware Resolution pane (Figure 7-27) and is related to these error types: `No hardware found` and `FICON EXP8S 10KM LX is not compatible with OSD`.

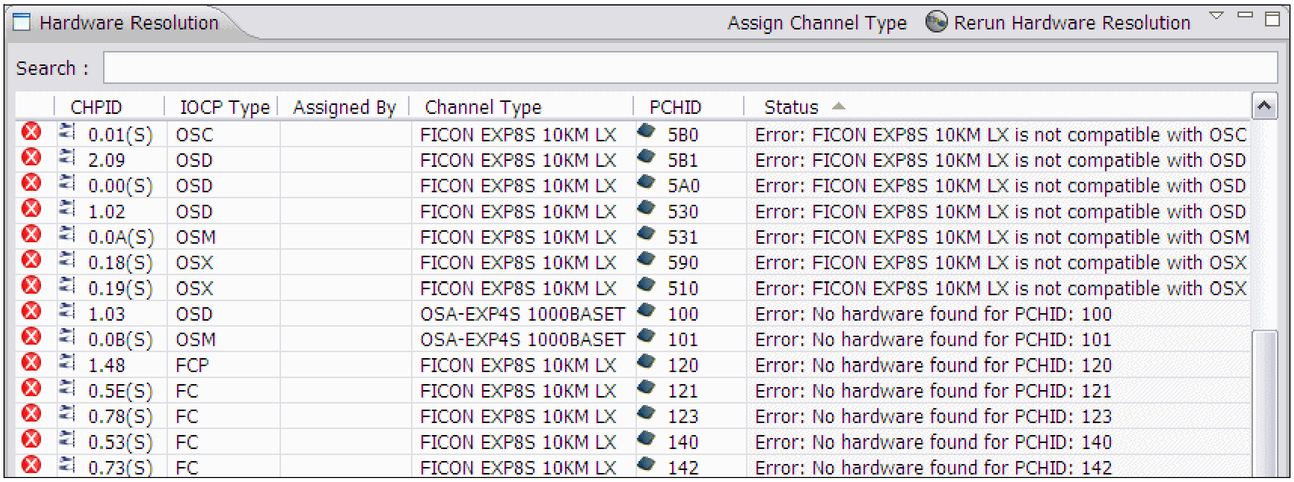

CHPID	IOCP Type	Assigned By	Channel Type	PCHID	Status
0.01(S)	OSC		FICON EXP8S 10KM LX	5B0	Error: FICON EXP8S 10KM LX is not compatible with OSC
2.09	OSD		FICON EXP8S 10KM LX	5B1	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.00(S)	OSD		FICON EXP8S 10KM LX	5A0	Error: FICON EXP8S 10KM LX is not compatible with OSD
1.02	OSD		FICON EXP8S 10KM LX	530	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.0A(S)	OSM		FICON EXP8S 10KM LX	531	Error: FICON EXP8S 10KM LX is not compatible with OSM
0.18(S)	OSX		FICON EXP8S 10KM LX	590	Error: FICON EXP8S 10KM LX is not compatible with OSX
0.19(S)	OSX		FICON EXP8S 10KM LX	510	Error: FICON EXP8S 10KM LX is not compatible with OSX
1.03	OSD		OSA-EXP4S 1000BASET	100	Error: No hardware found for PCHID: 100
0.0B(S)	OSM		OSA-EXP4S 1000BASET	101	Error: No hardware found for PCHID: 101
1.48	FCP		FICON EXP8S 10KM LX	120	Error: No hardware found for PCHID: 120
0.5E(S)	FC		FICON EXP8S 10KM LX	121	Error: No hardware found for PCHID: 121
0.78(S)	FC		FICON EXP8S 10KM LX	123	Error: No hardware found for PCHID: 123
0.53(S)	FC		FICON EXP8S 10KM LX	140	Error: No hardware found for PCHID: 140
0.73(S)	FC		FICON EXP8S 10KM LX	142	Error: No hardware found for PCHID: 142

Figure 7-27 CMT - Hardware resolution status errors

> **More information:** For more information about these error messages, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

The options that must be reset are as follows:

- Reset Incompatible (Hardware - I/O) Entries: (shown in example)
- Reset "Error: No hardware found" Entries: (shown in example)
- Reset "Select at least one adapter type": (shown in example)
- Reset "Required hardware for type IOCP_type not available": (*not* shown in example)
- Reset "PCHID_1 moved to new channel ID: PCHID_2": (*not* shown in example)

Reset Incompatible (Hardware - I/O) Entries

The Channel type that is assigned for the CHPID is not compatible with the IOCP type 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 PCHID. In the example, the IOCP type is OSD but the PCHID 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 (Figure 7-28 on page 381).

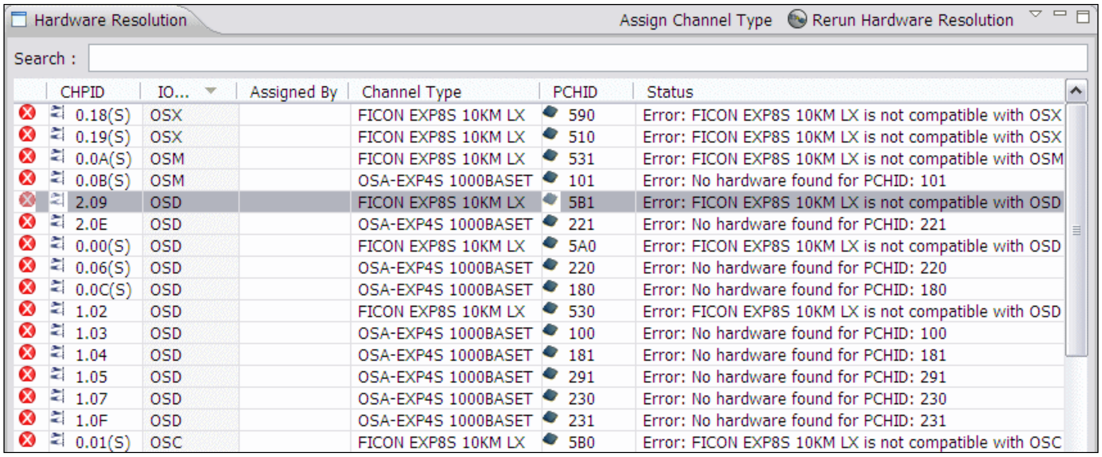

Figure 7-28 CMT - Channel_type is not compatible with IOCP_type

Select the channel type OSD. The Status is `Error: FICON EXP8S is not compatible with OSD`. Right-click in the row and select **Reset Incompatible (Hardware - I/O) Entries** to remove the PCHID values for only those rows (Figure 7-29).

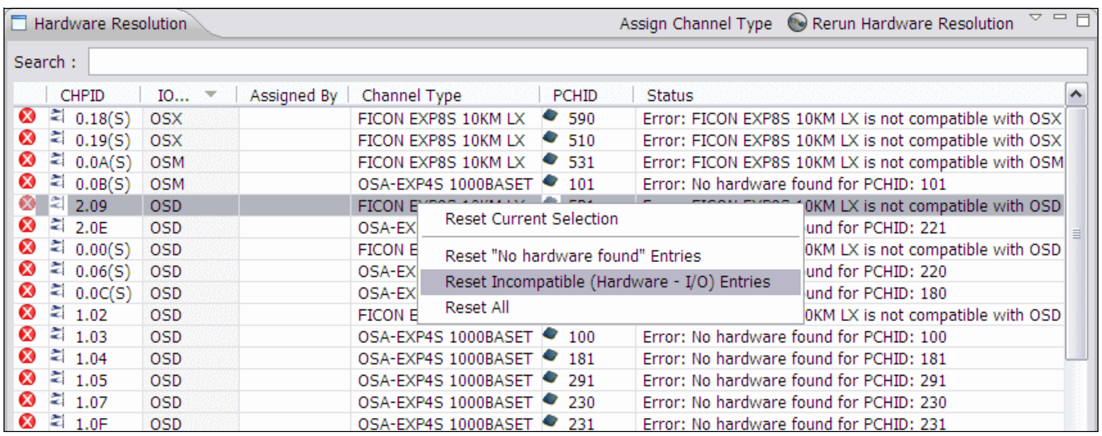

Figure 7-29 CMT - Channel_Type is not compatible with IOCP_type OSD

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 PCHIDs information (Figure 7-30).

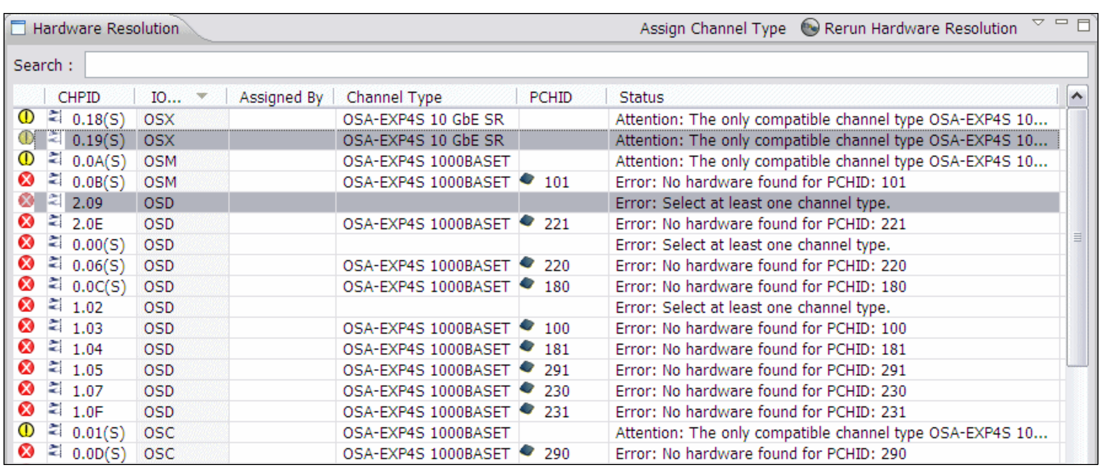

Figure 7-30 CMT - Results for reset of incompatible

The CHPID Mapping Tool now displays messages about any CHPID types that were imported from the IOCP input file (IODF) into the CMT that do not have any associated hardware support in the CFReport file (Figure 7-31). Click **OK**. The same figure also shows the Adapter Type Summary details.

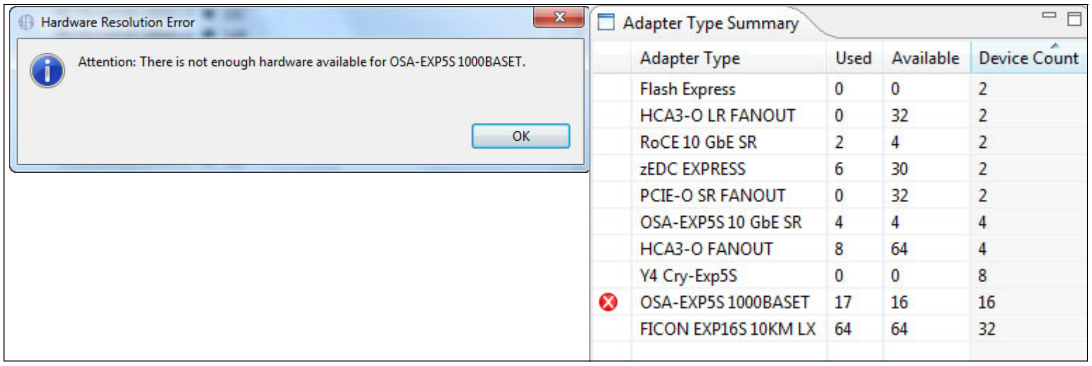

Figure 7-31 CMT - Required Hardware unavailable

Excessive numbers of *OSC* CHPID types are in the example IODF to show how the CHPID Mapping Tool handles this condition. For more information, see 7.2.3, “Overdefining channel paths on an XMP processor” on page 364.

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

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

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

To continue with this example, complete the following steps:

1. Return to the IODF and change the PCHID values for the OSD CHPIDs (or any other CHPIDs that have no 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 panel and selecting **Import IOCP File**.

Tip: If you look at the IOCP statements file now, the OSD CHPIDs have been omitted from the file, but are still defined in the IODF.

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

Reset “Error: No hardware found” Entries

An X in a red circle in the first column indicates an error, and the Status column provides the information with the value of `Error: No hardware found` (Figure 7-32).

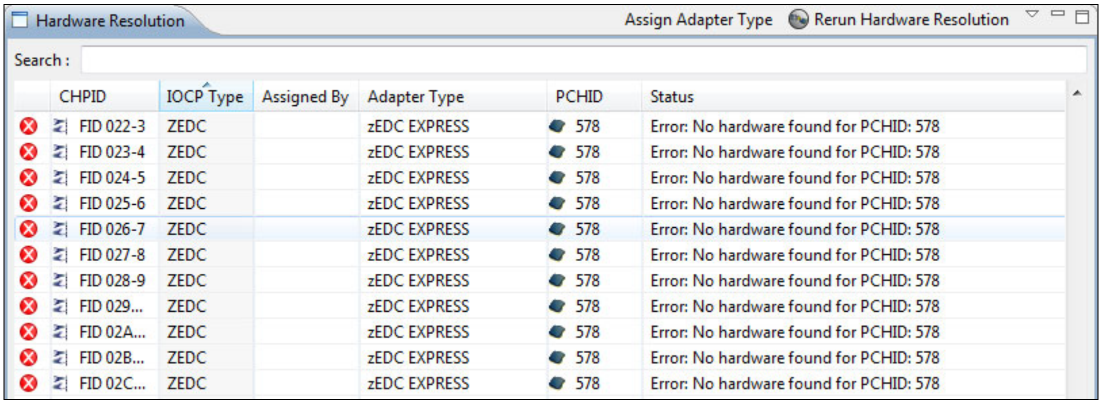

Figure 7-32 CMT - Error: No Hardware found

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 PCHID values for those rows (Figure 7-33).

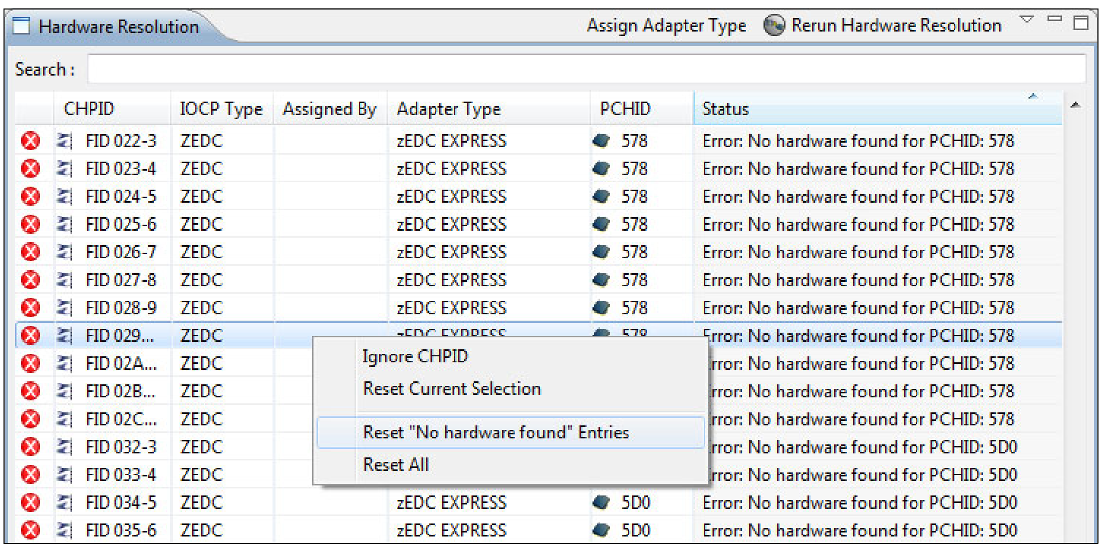

Figure 7-33 CMT - Resetting No Hardware found entries

The tool replaces the X with an *Attention* icon, changes the status message, and removes the PCHID information (Figure 7-34).

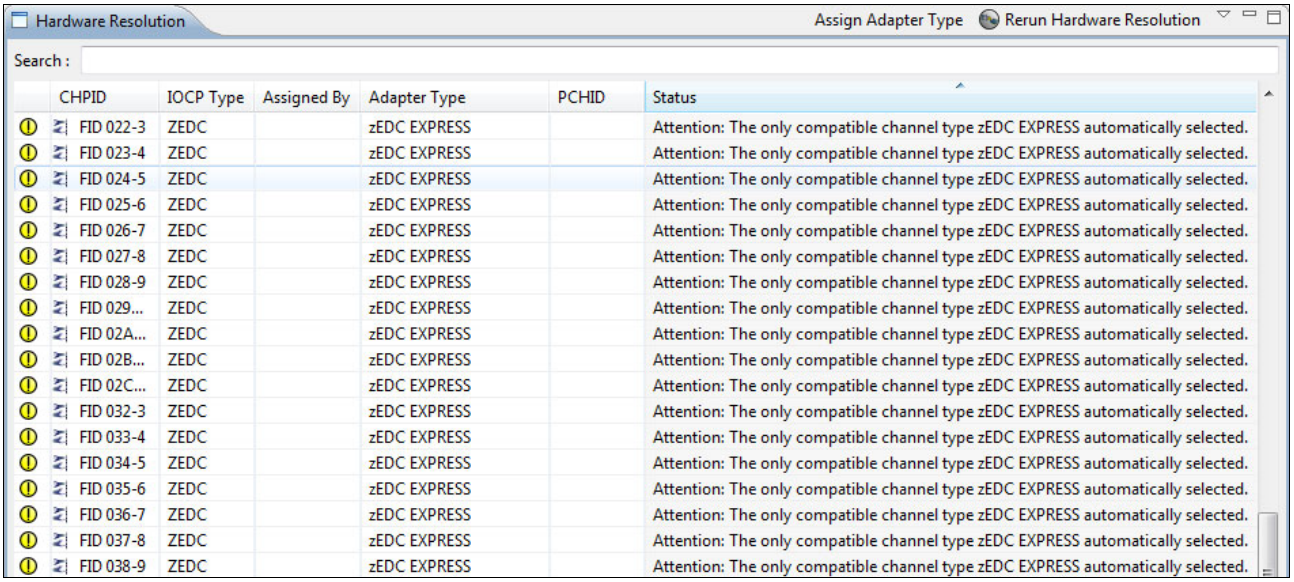

Figure 7-34 CMT - Results of resetting No hardware found

Reset "Select at least one adapter type"

The adapter type is not assigned to the current row. Assign an adapter type to the IOCP type:

1. Click the **Adapter Type** column in the target row. The tool displays an arrow in the Channel Type column of the target row (Figure 7-35).

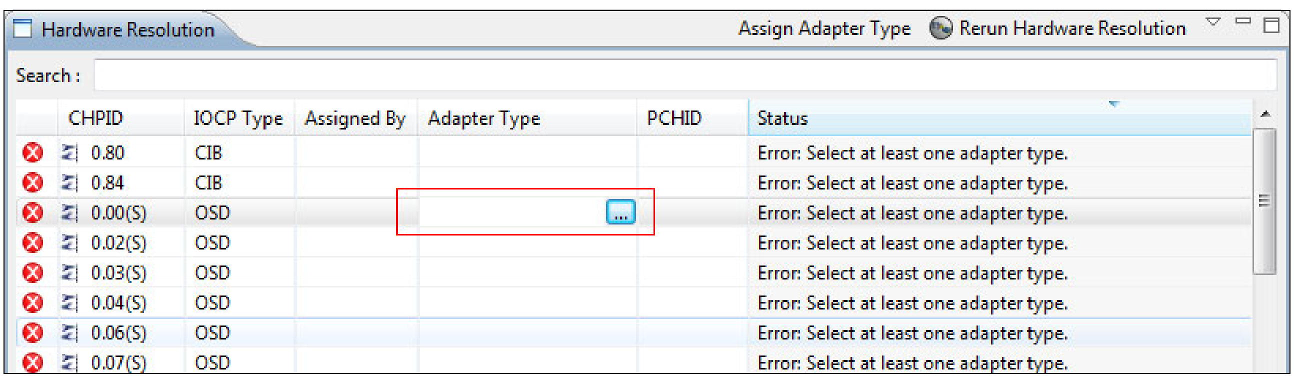

Figure 7-35 CMT - Selecting at least one adapter type

2. Click the ellipsis (...) box.

3. The tool displays a list of available and compatible card types for the CHPID as shown in Figure 7-36. Select an adapter type and click **OK**.

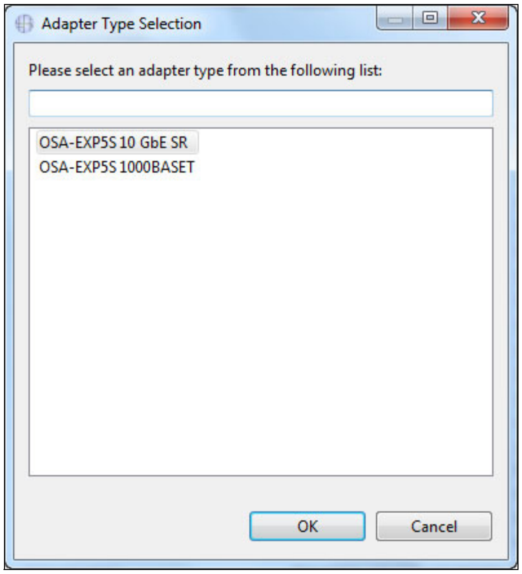

Figure 7-36 CMT - Adapter Type Selection

4. In the Adapter Type Summary tab, observe that the “Used” and “Available” totals change.

Reset “Required hardware for type IOCP_type not available”

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

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

You must change the IOCP file or obtain more hardware.

Reset “PCHID_1 moved to new channel ID: PCHID_2”

When moving from old hardware to new hardware, for example during a miscellaneous equipment specification (MES), the PCHID value assigned to a feature may change. This message indicates that the IOCP file contains a PCHID value for the old machine that is being removed. The PCHID value is changed from the old machine to the PCHID value for the new machine. For example, PCHID_1 is the first PCHID value representing the old hardware (for example, 1B0) and PCHID_2 is the new value representing the new hardware (for example, 533). In essence, the feature is present in both the old and new hardware, but its location (PCHID) has 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** button becomes available.

7.5.5 Manual mapping to resolve CIB CHPIDs

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

To resolve the CIB CHPIDs, assign the available CHPIDs. Click **Manual Mapping** (Figure 7-37).

Figure 7-37 CMT - Manual Mapping

Ensure that the tool is set to display Manual Mapping in the **Hardware -> I/O** format (Figure 7-38).

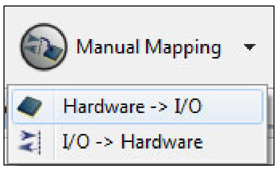

Figure 7-38 CMT - Manual Mapping of Hardware -> I/O

Click every row that has type HCA3-O in the Channel Type column. The tool displays all available CHPIDs with IOCP type (Figure 7-39).

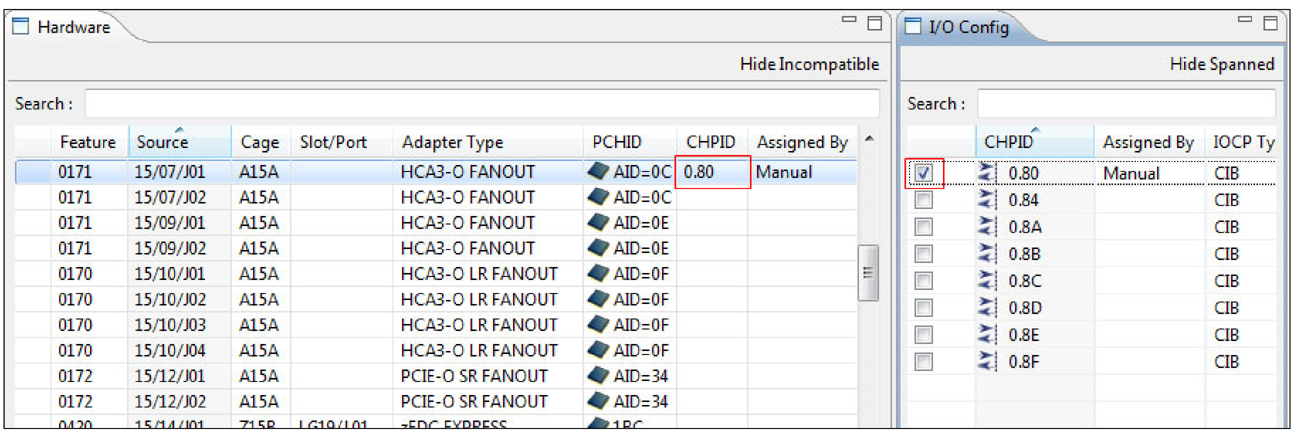

Figure 7-39 CMT - Adapter Type of HCA3 and associated CHPID assigned

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 HCA3-O adapter type, you see the `Multiple -->` value (Figure 7-40) inserted in the CHPID and Assigned By columns.

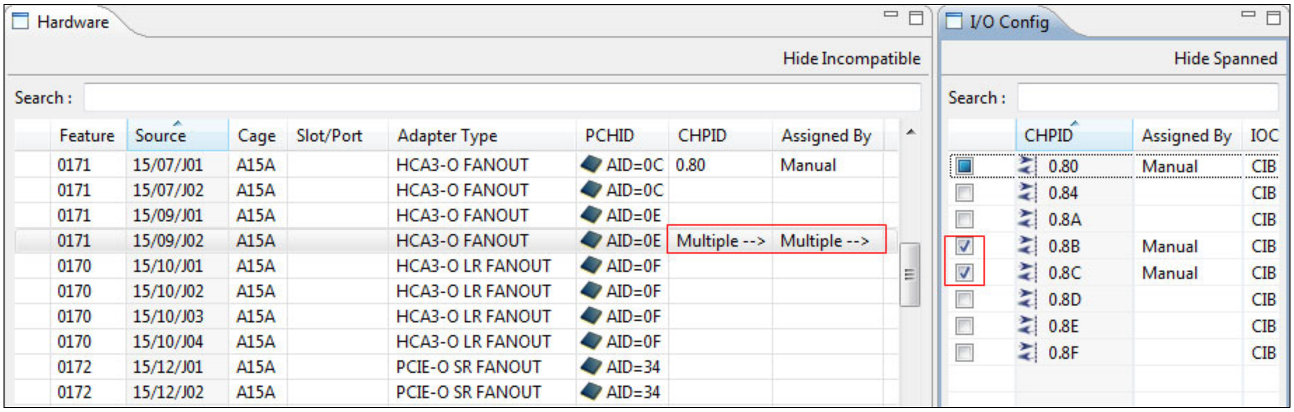

Figure 7-40 CMT - Adapter Type of HCA3 and associated multiple CHPID assigned

The **Automatic Mapping** button becomes available after you assign all the CHPIDs of IOCP type CIB.

7.5.6 Processing the CU Priority and Automatic Mapping

If you are importing an IOCP statements file from a 2817 or 2827 that had CU Priority values defined, review the CU Priority values first. The CHPID Mapping Tool can then perform the availability functions appropriately for a 2964.

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 these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices (Figure 7-41 on page 388). For the example, select the following two options and then click **OK**:
   - **Reset CHPIDs assigned by Automatic Mapping**
   - **Reset CHPIDs assigned by IOCP**

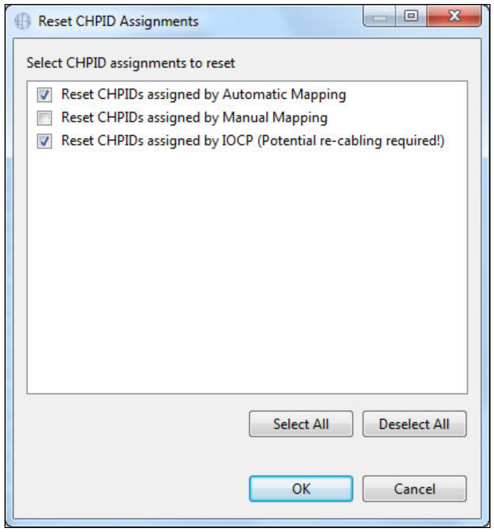

Figure 7-41 CMT - Reset CHPID Assignments

> **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 (Figure 7-42).

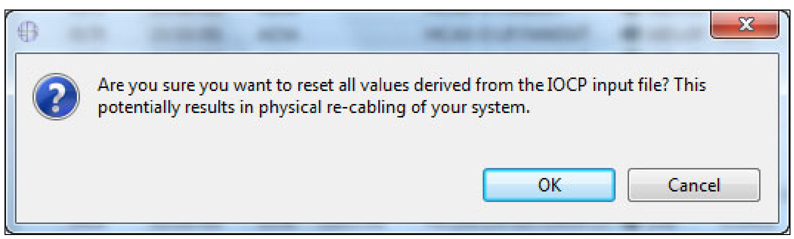

Figure 7-42 CMT - Reset CHPID assignments warning message

4. The 2964 has availability rules that differ from 2817 and 2827, so remove all PCHID 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 (Figure 7-43). Click **OK**.

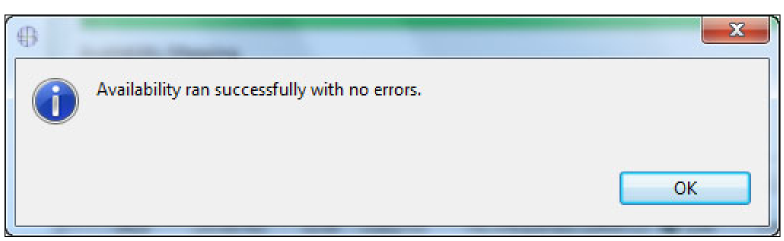

Figure 7-43 CMT - Availability ran successfully with no errors message

7. Click **OK** (Figure 7-44).

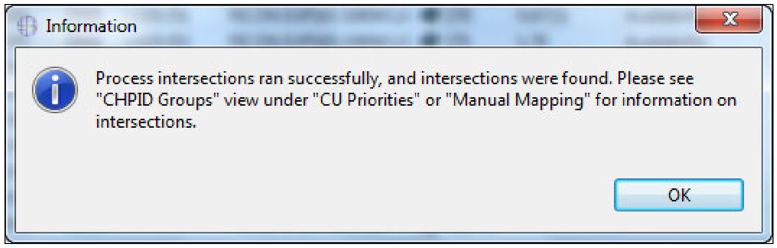

Figure 7-44 CMT - Process Intersections run successfully message

The possible intersects are as follows:

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 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 if they are acceptable (Figure 7-45). The example returned the “B” intersect.

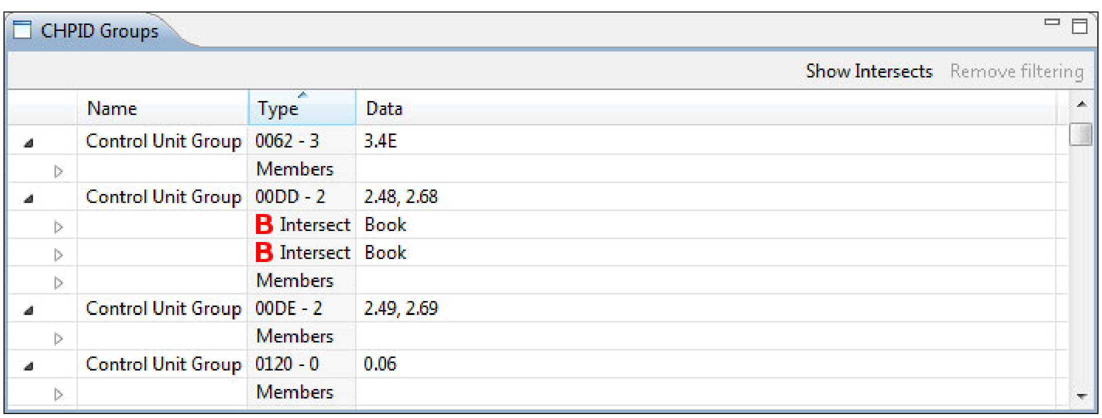

Figure 7-45 CMT - B Intersect examples

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.

Check and set values for items such as OSA-ICC CHPIDs and FCTC CHPIDs to ensure that the CHPID Mapping Tool allocates these CHPIDs with high PCHID availability.

1. Click **CU Priorities**. By default, this pane is in the center at the top.
2. In the CU Priorities pane, search in the CU Number column for the control units that you want to set a priority for.

3. Type 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. In the example, set the OSC type CU Numbers to priority 333 (Figure 7-46).

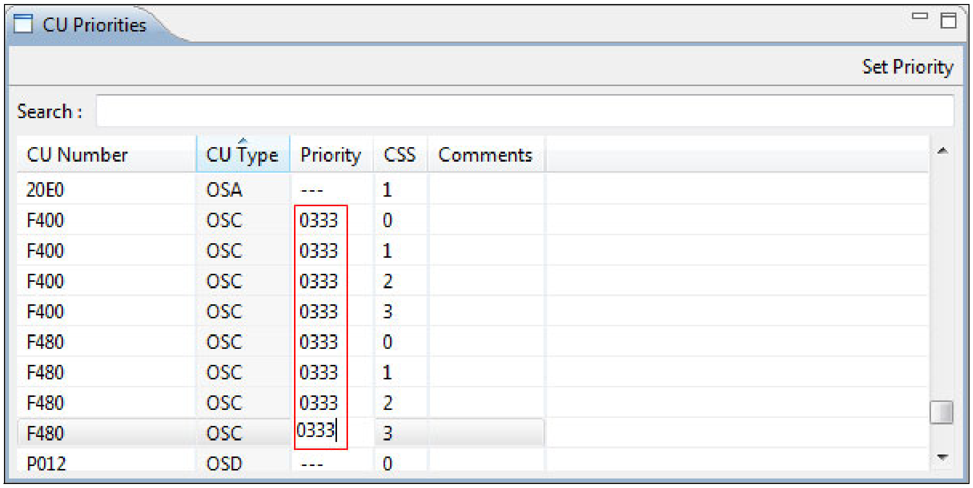

CU Number	CU Type	Priority	CSS	Comments
20E0	OSA	---	1	
F400	OSC	0333	0	
F400	OSC	0333	1	
F400	OSC	0333	2	
F400	OSC	0333	3	
F480	OSC	0333	0	
F480	OSC	0333	1	
F480	OSC	0333	2	
F480	OSC	0333	3	
P012	OSD	---	0	

Figure 7-46 CMT - Set CU Priorities

If coupling links are used by a CF image, group those links.

Group each set of CHPIDs going to a different CPC with a common priority. For example, suppose the CF image has four links (CHPIDs 40, 41, 42, and 43) and that 40 and 41 go to one CPC, and 42 and 43 go to a different CPC. In this case, give CHPIDs 40 and 41 one priority and CHPIDs 42 and 43 a different priority. The concept is the same regardless of the number of connecting CPCs or the number of links to each CPC.

Now perform the second availability function by completing these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices. Click **Reset CHPIDs assigned by Automatic Mapping**.
3. Click **OK**.

   The Hardware Resolution pane shows that the CHPID and Assigned By columns are no longer blank (Figure 7-47). The CMT assigned CHPIDs to PCHIDs and placed the Availability value in the Assigned By column, indicating that the CHPID values were assigned based on availability.

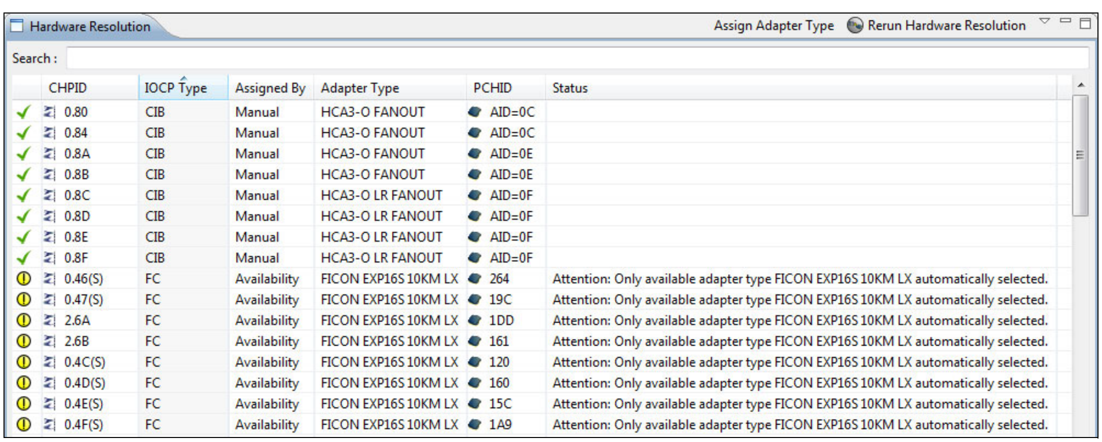

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
0.80	CIB	Manual	HCA3-O FANOUT	AID=0C	
0.84	CIB	Manual	HCA3-O FANOUT	AID=0C	
0.8A	CIB	Manual	HCA3-O FANOUT	AID=0E	
0.8B	CIB	Manual	HCA3-O FANOUT	AID=0E	
0.8C	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8D	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8E	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.8F	CIB	Manual	HCA3-O LR FANOUT	AID=0F	
0.46(S)	FC	Availability	FICON EXP16S 10KM LX	264	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.47(S)	FC	Availability	FICON EXP16S 10KM LX	19C	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
2.6A	FC	Availability	FICON EXP16S 10KM LX	1DD	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
2.6B	FC	Availability	FICON EXP16S 10KM LX	161	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4C(S)	FC	Availability	FICON EXP16S 10KM LX	120	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4D(S)	FC	Availability	FICON EXP16S 10KM LX	160	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4E(S)	FC	Availability	FICON EXP16S 10KM LX	15C	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.
0.4F(S)	FC	Availability	FICON EXP16S 10KM LX	1A9	Attention: Only available adapter type FICON EXP16S 10KM LX automatically selected.

Figure 7-47 CMT - CHPIDs assigned

The possible Assigned By column values are as follows:

Manual You made the assignment by using manual mapping.

Automatic You made the assignment by using automatic mapping.

IOCP The IOCP source made the assignment.

Config File The CHPID Mapping Tool forced an assignment because of configuration file requirements.

> **More information:** See the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

4. You can now display the results of the channel mapping. You can also sort the report in various ways. For example, to see how the CHPID Mapping Tool ranked the control units, select the **CU Priorities** pane and click the **Priority** column (Figure 7-48).

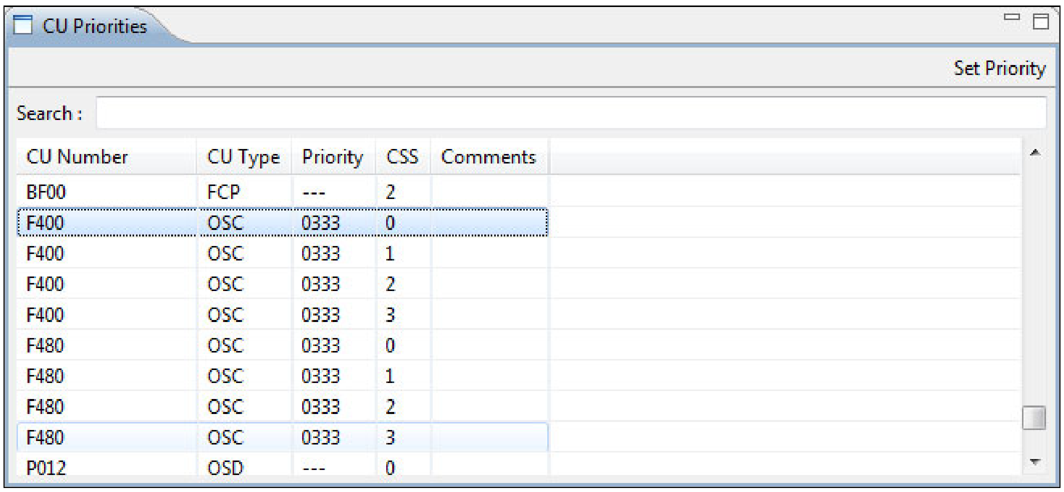

CU Number	CU Type	Priority	CSS	Comments
BF00	FCP	---	2	
F400	OSC	0333	0	
F400	OSC	0333	1	
F400	OSC	0333	2	
F400	OSC	0333	3	
F480	OSC	0333	0	
F480	OSC	0333	1	
F480	OSC	0333	2	
F480	OSC	0333	3	
P012	OSD	---	0	

Figure 7-48 CMT - CU Priorities showing assigned priorities

The example illustrates how CU Priority values are represented in the IOCP file.

> **Tip:** The control unit priorities are stored in the IOCP output file created by the CMT that gets migrated back into HCD. HCD maintains these priorities and outputs them when it creates another IOCP deck. They are in the form of commented lines at the end of the IOCP deck, as shown here:

```
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
```

7.5.7 CHPIDs not connected to control units

In the CU Priorities window, click in the **CU Number** column (Figure 7-49). The CHPID Mapping Tool shows, at the end of the list, all CHPIDs 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.

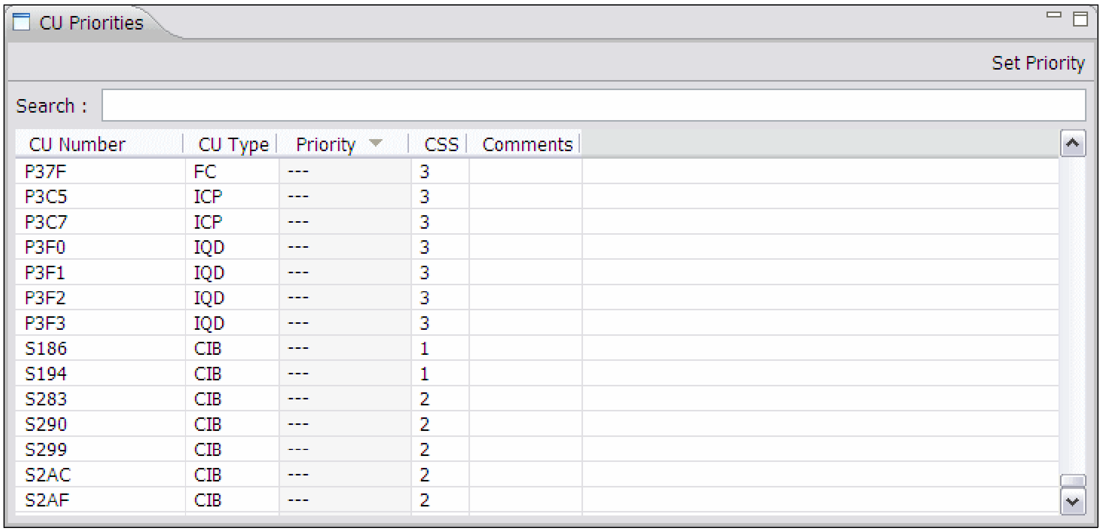

CU Number	CU Type	Priority	CSS	Comments
P37F	FC	---	3	
P3C5	ICP	---	3	
P3C7	ICP	---	3	
P3F0	IQD	---	3	
P3F1	IQD	---	3	
P3F2	IQD	---	3	
P3F3	IQD	---	3	
S186	CIB	---	1	
S194	CIB	---	1	
S283	CIB	---	2	
S290	CIB	---	2	
S299	CIB	---	2	
S2AC	CIB	---	2	
S2AF	CIB	---	2	

Figure 7-49 CMT - 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 need to 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 there are extra CHPIDs for anticipated new control units, you might want to group these CHPIDs with a common priority. Having a common priority allows the availability mapping function to pick PCHIDs that can afford your new control unit availability.

7.5.8 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**. Figure 7-50 shows the options to create a Preview Report or Save Report.

Figure 7-50 CMT - 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 (Figure 7-51).

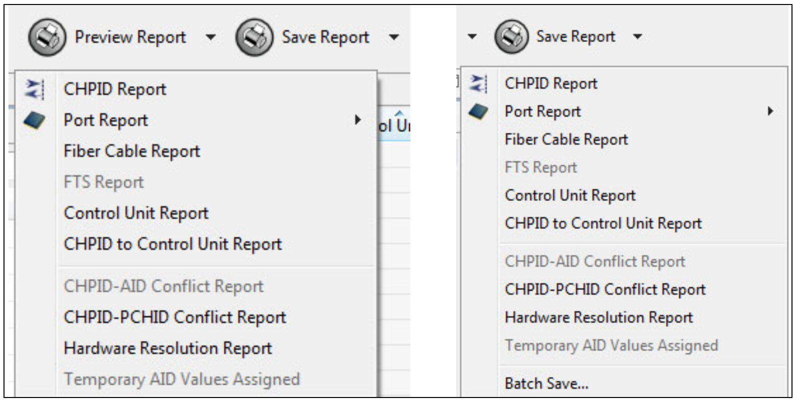

Figure 7-51 CMT - Preview Report and Save Report menus

For simplicity, only three reports are described in this example: the CHPID Report; the Port Report, sorted by location; and the 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 PCHID or cage/slot/port information before system delivery.

CHPID Report

To create the CHPID report, complete the following steps:

1. Click **Preview Report** → **CHPID Report** (Figure 7-52).

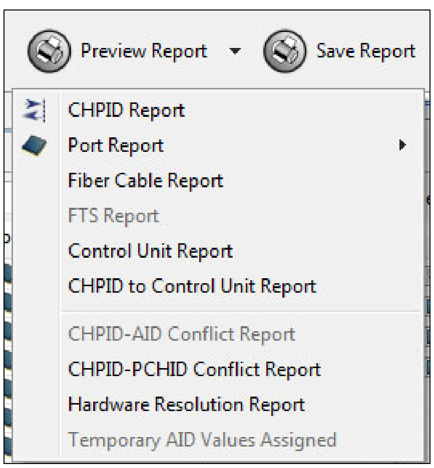

Figure 7-52 CMT - Preview report - CHPID Report

The CHPID Mapping Tool displays the CHPID Report in a Report tab within the CMT (Figure 7-53).

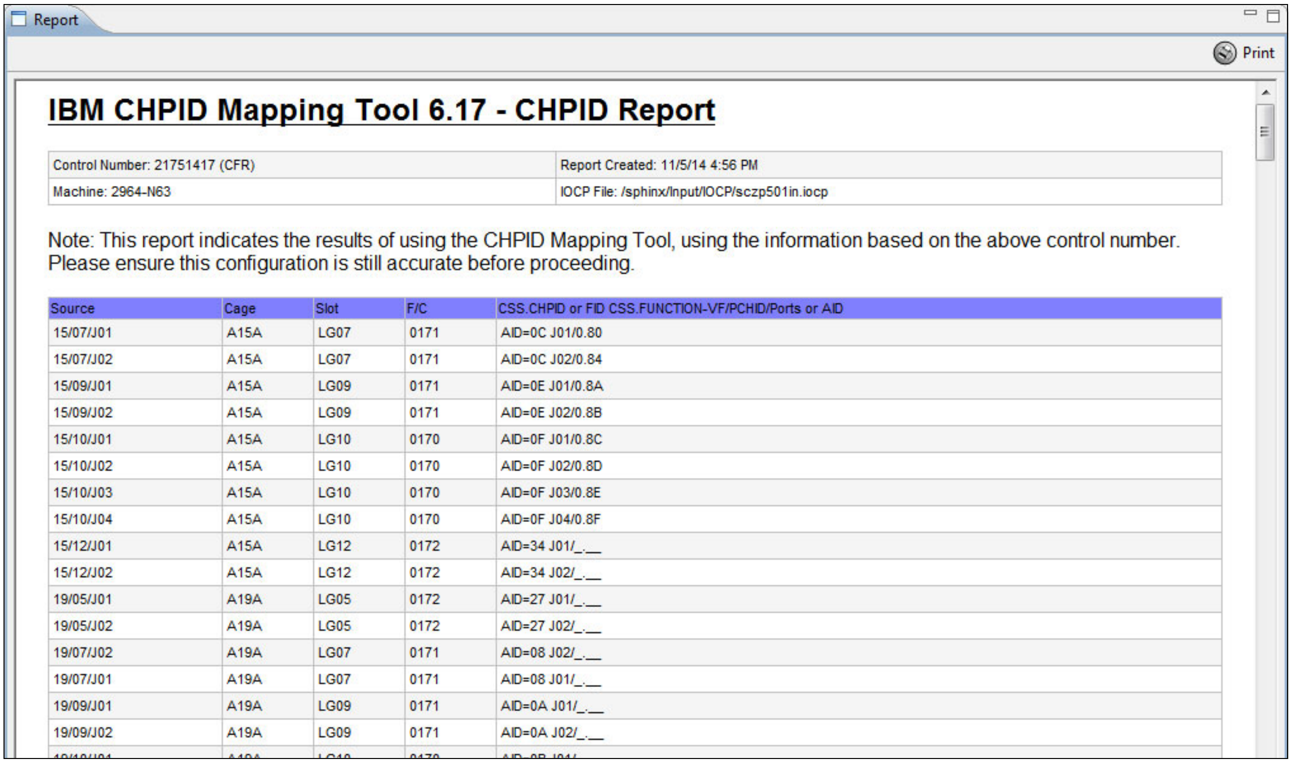

Report

Print

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 4:56 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__

Figure 7-53 CMT - 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 (Figure 7-54). 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**.

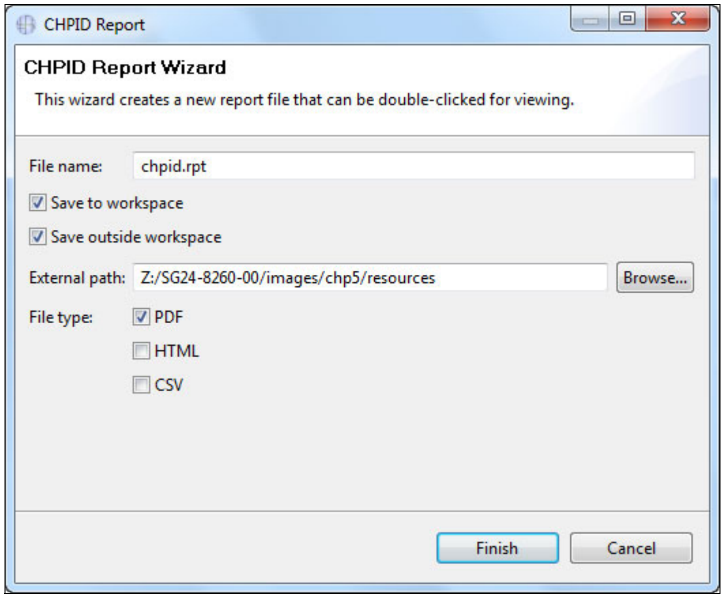

Figure 7-54 CMT - Save CHPID Report

The CHPID Report is created by the CHPID Mapping Tool (Figure 7-55).

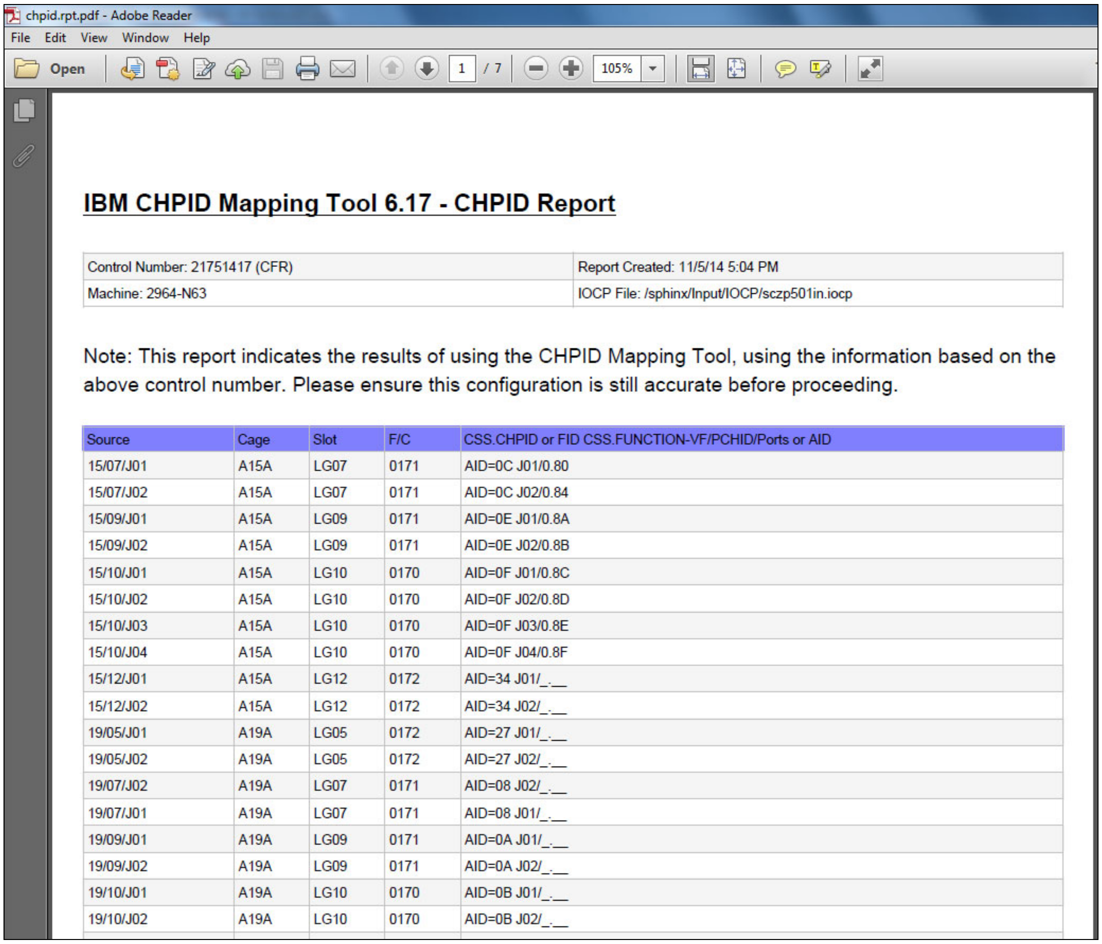

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:04 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__
19/10/J01	A19A	LG10	0170	AID=0B J01/_.__
19/10/J02	A19A	LG10	0170	AID=0B J02/_.__

Figure 7-55 CMT - CHPID Report example in PDF format

At the end of this CHPID Report is a list of CHPIDs with modified PCHID/AID assignments (Figure 7-56). This report is valuable for moving cables.

List of CHPIDs having modified PCHID/AID assignments

Note: For CHPIDs that had PCHID/AID assignments in the IOCP file that was loaded for this session of the Mapping Tool.

CHPIDs or FUNCTIONs	Previous PCHID/AID-Port	PCHID/AID-Port	Current Location	F/C
0.80	09-1	0C-1	A15ALG07J.01	0171
0.84	09-1	0C-02	A15ALG07J.02	0171
0.8A	09-1	0E-1	A15ALG09J.01	0171
0.8B	09-1	0E-02	A15ALG09J.02	0171
0.8C	09-1	0F-01	A15ALG10J.01	0170
0.8D	09-1	0F-02	A15ALG10J.02	0170
0.8E	09-1	0F-03	A15ALG10J.03	0170
0.8F	09-1	0F-04	A15ALG10J.04	0170
FID 000	544	208	Z08BLG03J.01	0411
FID 010	5EC	140	Z22BLG20J.01	0411
FID 020-1	578	1BC	Z15BLG19J.01	0420
FID 021-2	578	27C	Z08BLG38J.01	0420
FID 022-3	578	Not Assigned		
FID 023-4	578	Not Assigned		

Figure 7-56 CMT - List of CHPIDs that have modified PCHID/AID assignments

Port Report, sorted by location

To create the Port Report, 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 (Figure 7-57).

IBM CHPID Mapping Tool 6.17 - CHPID to Port Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:12 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
A15A	07/LG	AID=0C / J.01	15/07	HCA3-O FANOUT	0.80	Manual
A15A	07/LG	AID=0C / J.02	15/07	HCA3-O FANOUT	0.84	Manual
A15A	09/LG	AID=0E / J.01	15/09	HCA3-O FANOUT	0.8A	Manual
A15A	09/LG	AID=0E / J.02	15/09	HCA3-O FANOUT	0.8B	Manual
A15A	10/LG	AID=0F / J.01	15/10	HCA3-O LR FANOUT	0.8C	Manual
A15A	10/LG	AID=0F / J.02	15/10	HCA3-O LR FANOUT	0.8D	Manual
A15A	10/LG	AID=0F / J.03	15/10	HCA3-O LR FANOUT	0.8E	Manual
A15A	10/LG	AID=0F / J.04	15/10	HCA3-O LR FANOUT	0.8F	Manual
A15A	12/LG	AID=34 / J.01	15/12	PCIE-O SR FANOUT		
A15A	12/LG	AID=34 / J.02	15/12	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.01	19/05	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.02	19/05	PCIE-O SR FANOUT		
A19A	07/LG	AID=08 / J.01	19/07	HCA3-O FANOUT		
A19A	07/LG	AID=08 / J.02	19/07	HCA3-O FANOUT		
A19A	09/LG	AID=0A / J.01	19/09	HCA3-O FANOUT		

Figure 7-57 CMT - CHPID to Port Report, sorted by location

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 (Figure 7-58).

IBM CHPID Mapping Tool 6.17 - CHPID to CU Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:15 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
0	00	OSD	19/12/J01	Z08B LG07 J00J01	214	2040	OSA	---
0	01	OSC	19/03/J01	Z15B LG21 J00J01	1C4	F400	OSC	0333
0	02	OSD	15/03/J01	Z15B LG35 J00J01	1F0	2080	OSA	---
0	03	OSD	15/15/J01	Z22B LG17 J00J01	134	20A0	OSA	---
0	04	OSD	15/02/J01	Z22B LG36 J00J01	174	20C0	OSA	---
0	06	OSD	19/02/J01	Z22B LG21 J00J01	144	0120	OSA	---
0	07	OSD	15/02/J01	Z22B LG37 J00J01	178	0140	OSA	---
0	0C	OSD	19/03/J01	Z15B LG22 J00J01	1C8	2060	OSA	---
0	0D	OSC	15/15/J01	Z22B LG18 J00J01	138	F480	OSC	0333
0	12	OSD	15/05/J01	Z08B LG36 J00J01	274	P012	OSD	---
0	13	OSD	19/14/J01	Z15B LG06 J00J01	190	P013	OSD	---
0	40	FC	15/14/J01	Z15B LG12 D1	1A4	1000	2107	---
						1200	2107	---
						1400	2107	---

Figure 7-58 CMT - CHPID to CU Report

7.5.9 Creating an updated IOCP

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

> **Using HCM:** You might prefer to use HCM to transfer the updated IOCP statements file back to the host. However, first run the next step in the CHPID Mapping Tool to create the updated IOCP file.

To create the IOCP, complete the following steps:

1. Select **File** → **Export IOCP input file** (Figure 7-59).

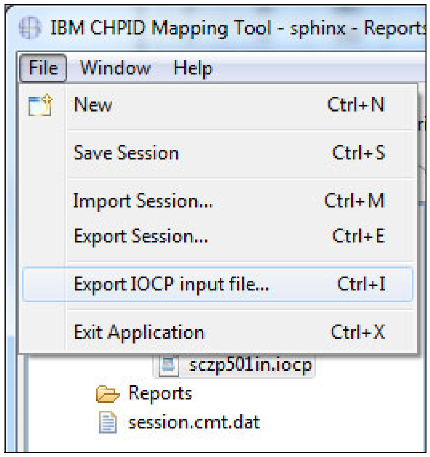

Figure 7-59 CMT - Export IOCP input file

2. Enter the Export Path and IOCP Name for the IOCP output file and click **Finish** (Figure 7-60).

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

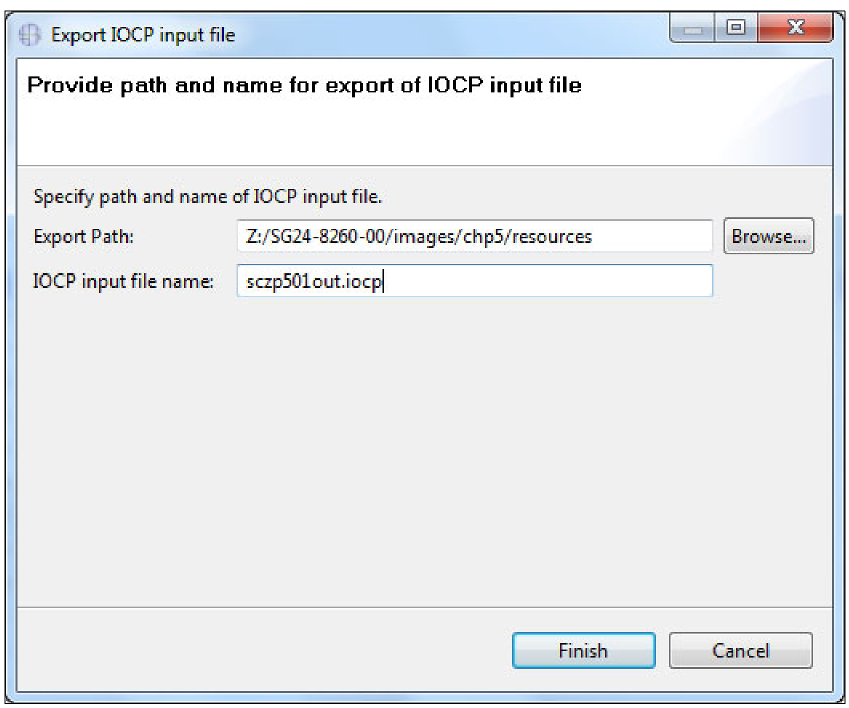

Figure 7-60 CMT - Export IOCP File

3. Select **File** → **Save Session** (Figure 7-61).

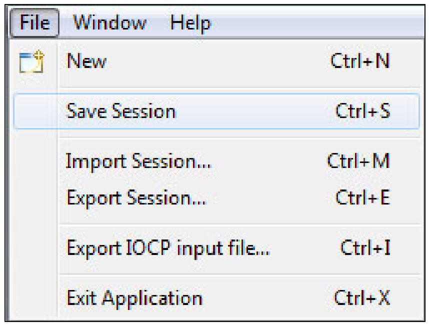

Figure 7-61 CMT - Save session

You might want to save your project before exiting the CMT application.

7.6 HCD: Updating the 2964 work IODF with PCHIDs

After you map the PCHIDs to CHPIDs by using the CHPID Mapping Tool, transfer this information back into HCD. To update the IODF with the PCHIDs, follow these following steps:

1. Upload the IOCP file that was created by the CMT (`SCZP501out.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 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. They are in the form of commented lines at the end of the IOCP deck (Example 7-2).

Example 7-2 HCD - Updated IOCP statements file with CMT statements

```
         IODEVICE ADDRESS=(FF3F,007),CUNUMBR=(FFFB),UNIT=CFP
         IODEVICE ADDRESS=(FF58,007),CUNUMBR=(FFFB),UNIT=CFP
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 0061.0=0400,0062.0=0400,0B00.0=0500,0B01.0=0500,0B20.0=0520
*CMT* 1000.0=0100,1100.0=0100,1200.0=0100,1300.0=0100,1400.0=0100
*CMT* 1500.0=0100,1600.0=0100,1700.0=0100,2100.0=0320,2120.0=0320
*CMT* 4B00.0=0600,4B08.0=0600,5B00.0=0600,5B08.0=0600,9000.0=0120
*CMT* 9080.0=0120,9100.0=0120,9180.0=0120,9200.0=0120,9280.0=0120
*CMT* 9300.0=0120,9380.0=0120,9400.0=0120,9480.0=0120,9500.0=0120
*CMT* 9580.0=0120,9600.0=0120,9680.0=0120,9800.0=0120,9880.0=0120
*CMT* 9900.0=0120,9980.0=0120,9A00.0=0120,9A80.0=0120,9B00.0=0120
*CMT* 9B80.0=0120,9C00.0=0120,9C80.0=0120,9D00.0=0120,9D80.0=0120
*CMT* 9E00.0=0120,9E80.0=0120,A000.0=0100,A100.0=0100,A200.0=0100
*CMT* A300.0=0100,A400.0=0100,A500.0=0100,A600.0=0100,A700.0=0100
*CMT* C000.0=0160,C100.0=0160,C200.0=0160,C300.0=0160,C400.0=0160
*CMT* C500.0=0160,C600.0=0160,C700.0=0160,C800.0=0160,C900.0=0160
*CMT* CA00.0=0160,CB00.0=0160,CC00.0=0160,CD00.0=0160,CE00.0=0160
*CMT* CF00.0=0160,D000.0=0180,D100.0=0180,D200.0=0180,D300.0=0180
*CMT* D400.0=0180,D500.0=0180,D800.0=0180,D900.0=0180,DA00.0=0180
*CMT* DB00.0=0180,DC00.0=0180,DD00.0=0180,DE00.0=0180,DF00.0=0180
*CMT* F401.0=0300,F481.0=0300,PF030.0=0001,PF031.0=0001,PF020.0=0001
*CMT* PF032.0=0001,PF021.0=0001,PF010.0=0001,PF022.0=0001,PF011.0=0001
*CMT* PF000.0=0001
******************************** Bottom of Data ************************
```

Important: Do not edit the CMT comments that are contained in the IOCP output file manually. If priorities must be changed, use the CMT.

2. On the HCD main panel, enter the work IODF name that you used to create the IOCP input data set for the CHPID Mapping Tool. Select option **5. Migrate configuration data** (Figure 7-62).

```
                          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 . . . 'SYS6.IODFB7.WORK'                    +
```

Figure 7-62 HCD - Hardware Configuration: migrate configuration data

3. From the Migrate Configuration Data panel (Figure 7-63), 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
```

Figure 7-63 HCD - Migrate Configuration Data: migrate IOCP/OS data.

4. The Migrate IOCP Data panel opens (Figure 7-64). 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 OS configuration that is associated with the processor.

IOCP only input data set This data set was specified when the `sczp501out.iocp` file was uploaded to z/OS.

Processing mode Select option **2** to save the results of the migration. (Before using option 2, try to migrate 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 . . . . . . . . . . . . SCZP501   +   CSS ID . . . . . . _  +
OS configuration ID  . . . . . . . . L06RMVS1  +

Combined IOCP/MVSCP input data set . _____________________________________
IOCP only input data set . . . . . . 'SYS6.IODFB7.IOCPOUT.SCZP501'
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)
```

Figure 7-64 HCD - 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 (Figure 7-65 on page 403).

```
------------------------- 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
#                     SYS6.IODFB7.WORK.
***************************** Bottom of data ******************************
```

Figure 7-65 HCD - Migration Message List: successful message

The work IODF now contains both the CHPID definitions and the mapping to PCHIDs 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.

7.7 HCD: Building the 2964 production IODF

To use the definitions that were updated in HCD, create a 2964 production IODF from your work IODF. Then, remotely or locally write the IODF to the 2964 IOCDS by using Write IOCDS in preparation for the upgrade.

To create a production IODF, complete the following steps:

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 7-66).

```
                           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.IODFB7.WORK'                     +
```

Figure 7-66 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 7-67). 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
```

Figure 7-67 HCD - Activate or Process Configuration Data: build production IODF

3. HCD displays the Message List panel (Figure 7-68). Verify that you have only severity "W" (warning) messages and that they are normal for your configuration. Correct any other messages that should not occur and try to build the production IODF again. Continue this process until you have no messages that indicate problems.

```
------------------------------ Message List -----------------------------
  Save  Query  Help
---------------------------------------------------------------------------
                                                                Row 1 of 173
Command ===> ___________________________________________ Scroll ===> CSR

Messages are sorted by severity. Select one or more, then press Enter.
/ Sev Msg. ID  Message Text
_ W   CBDG159I Switch control unit(s) 0061 and device(s) 0061 defined,
#              but not yet connected to both a processor and an
#              operating system.
_ W   CBDG159I Switch control unit(s) 0062 and device(s) 0062 defined,
#              but not yet connected to both a processor and an
#              operating system.
_ W   CBDG159I Switch control unit(s) 0063 and device(s) 0063 defined,
#              but not yet connected to both a processor and an
#              operating system.
```

Figure 7-68 HCD - Message List: building production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (Figure 7-69). Complete the Production IODF name and Volume serial number fields, and then press Enter.

```
------------- Build Production I/O Definition File ---------


Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODFB7.WORK'

Production IODF name . 'SYS6.IODFB7'
Volume serial number . IODFPK  +

Continue using as current IODF:
2   1.  The work IODF in use at present
    2.  The new production IODF specified above
```

Figure 7-69 HCD - Build Production I/O Definition File: data fields to be updated

6. The Define Descriptor Fields panel opens (Figure 7-70). 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.IODFB7'

Descriptor field 1  . . . SYS6
Descriptor field 2  . . . IODFB7
```

Figure 7-70 HCD - 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.IODFB7 created.`

7.8 HCD/HMC: Loading the 2964 processor IOCDS

You now have a production IODF, named `SYS6.IODFB7`. Now the IOCDS component of the IODF must be updated on the *replacement* CPC that is being installed (for example, SCZP501) and activated (POR) using this IOCDS.

The final step is to perform an initial program load (IPL) of the processor using this IODF. (Describing how to perform the IPL of the new hardware is beyond the scope of this book.)

The two possible ways to load the IOCP Statements onto the 2964 Support Element IOCDS are as follows:

- HCD, using option 2.11
- The HMC/SE, by using the Stand-Alone Input/Output Configuration Program

Although both are valid methods to write the new configuration to the IOCDS, using HCD option 2.11 is the preferred method. However, your new 2964 processor and Support Element might not be connected to the system where the configuration was generated or cannot be connected to any system where HCD is running. In that case, use the Stand-Alone IOCP process.

7.8.1 Updating the IOCDS using HCD option 2.11

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

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 7-71). Ensure that the IODF is the production IODF that was created in 7.7, “HCD: Building the 2964 production IODF” on page 403 and then press Enter.

```
                           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.IODFB7'                          +
```

Figure 7-71 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 7-72).
   Select option **11. Build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes**.

```
----- 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
```

Figure 7-72   HCD - Activate or Process Configuration data: build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes

Consideration: This example assumes that you have connectivity to the new 2964 over the HMC LAN to create an IOCDS from which you do a power-on reset.

If the new 2964 is not accessible from the HMC LAN, copy the IOCP statements onto a USB flash memory drive. You can then import them into the 2964 HMC to run a stand-alone IOCP. Creating a file on a USB flash memory drive can be done using the same process that is used to create an IOCP input file for the CHPID Mapping Tool.

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

3. The System z Cluster List panel opens (Figure 7-73). In the list, use a forward slash (/) to select the new 2964 to update one of its IOCDSs. Then, press Enter.

```
                            System z Cluster List                  Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
/ USIBMSC.SCZP501    2964   N63     SCZP501
******************************* Bottom of data ********************************
```

Figure 7-73   HCD - System z Cluster List: selecting processor for IOCDS update

4. The Actions on selected CPCs panel opens (Figure 7-74). 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)
```

Figure 7-74 HCD - Actions on selected CPCs: work with IOCDSs

5. The IOCDS List panel opens (Figure 7-75). Select the IOCDS that you want to update for the 2964 installation by typing a forward slash (/) next to it, and 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.SCZP501  IODFB6   LPAR   Alternate  No        No          No
_ A1.SCZP501  IODF40   LPAR   Alternate  No        No          No
_ A2.SCZP501  IODF42   LPAR   Alternate  No        No          No
_ A3.SCZP501           LPAR   POR        Yes       No          Yes-POR
****************************** Bottom of data ********************************
```

Figure 7-75 HCD - IOCDS List: selecting IOCDS for update

6. The Actions on selected IOCDSs panel opens (Figure 7-76). Select option **1. Update IOCDS** and press Enter.

```
--------------- Actions on selected IOCDSs ------


Select by number or action code and press Enter.

1_  1.  Update IOCDS . . . . . . . . . . . . (u)
    2.  Switch IOCDS . . . . . . . . . . . . (s)
    3.  Enable write protection  . . . . . . (e)
    4.  Disable write protection . . . . . . (w)
```

Figure 7-76 HCD - Actions on selected IOCDSs: update IOCDS

7. The Build IOCDSs panel opens (Figure 7-77). Verify that all the information is correct. Complete the Title1 field and press Enter.

```
------------------------------ Build IOCDSs -------------------------------
                                                                Row 1 of 1
Command ===> ___________________________________________ Scroll ===> CSR

Specify or revise the following values.

IODF name . . . . . . . . . : 'SYS6.IODFB7'

Title1 . IODFB7_________________________________________________________
Title2 : SYS6.IODFB7 - 2014-11-10 13:35

                           Write IOCDS in
IOCDS        Switch IOCDS  preparation of upgrade
A0.SCZP501   No            No
**************************** Bottom of data ******************************
```

Figure 7-77 HCD - Build IOCDSs: verifying IODF

8. The Job Statement Information panel opens (Figure 7-78). Complete the job statements as required by the installation and press Enter. HCD submits the job to update the IOCDS.

```
------------------------ Job Statement Information -------------------


Specify or revise the job statement information.


Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 7-78 HCD - Job Statement Information: option to override job statement cards

9. Verify the job output to ensure that the IOCDS was written without error and to the correct IOCDS. You receive a message similar to the following message:

```
ICP057I IOCP JOB WIOCP    SUCCESSFUL.  LEVEL A0 IOCDS REPLACED.

Sev  Msgid    Message Text
 I  CBDA674I  IOCP successfully completed for A0.SCZP501.
```

10. Now iff you return to HCD option 2.11 and view the IOCDS, the SNA Address is at `USIBMSC.SCZP501` (Figure 7-79).

```
                             System z Cluster List                   Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
s USIBMSC.SCZP501    2964   N63     SCZP501
****************************** Bottom of data *******************************
```

Figure 7-79 HCD - System z Cluster List: selecting processor for IOCDS verify

Figure 7-80 shows the updated IOCDS with Alternate status.

```
                                   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.SCZP501   IODFB7   LPAR   Alternate   No        Yes         No
_ A1.SCZP501   IODF40   LPAR   Alternate   No        No          No
_ A2.SCZP501   IODF42   LPAR   Alternate   No        No          No
_ A3.SCZP501   IODF43   LPAR   POR         Yes       No          Yes-POR
****************************** Bottom of data *******************************
```

Figure 7-80 HCD - IOCDS List: IOCDS verified

7.8.2 Updating the IOCDS using Stand-Alone Input/Output Config Program

Copy the IOCP statements that were generated by using HCD option **2.3. Build IOCP input data set** onto a USB flash memory drive and retain it.

Tip: For more information, see the *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152.

To update the IOCDS, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, as opposed to a remote web browser. Select the new 2964, assuming it was defined to the Defined CPCs Work Area.
2. Perform a power-on reset using one of the Reset profiles and Starter IOCDSs provided with the processor during installation. This action creates an environment on the processor in which you can run the Stand-Alone IOCP process.
3. When the power-on reset is complete, activate one of the logical partitions with at least 128 MB of storage. Use this partition to run the I/O Configuration Program.
4. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
5. Under Systems or Members, click the system to select it (in this example, SCZP501).

6. On the Tasks tab, click **Recovery** → **Single Object Operations** → **Yes**.
7. Under Systems Management, click the system to select it (in this example, SCZP501).
8. Under Partitions, select the LPAR you want to use to run the Stand-Alone IOCP program (in this example, A0B) as shown in Figure 7-81.

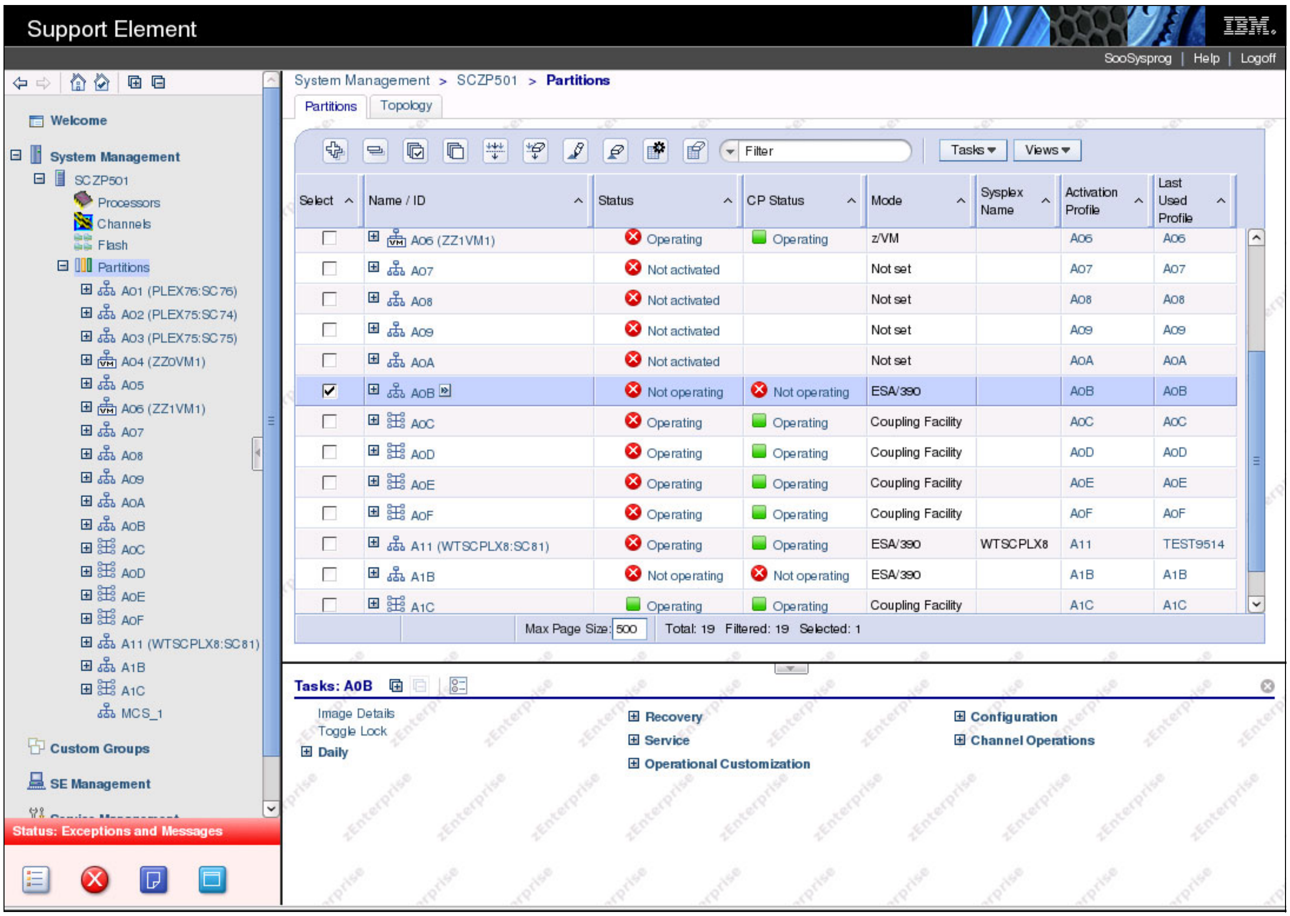

Figure 7-81 HMC - SAIOCP: partition selected for SAIOCP program load

9. On the Tasks tab, click **Configuration** → **Input/output (I/O) Configuration**.
10. Select the data set into which you want to import the IOCDS (in this example, A0) as shown in Figure 7-82.

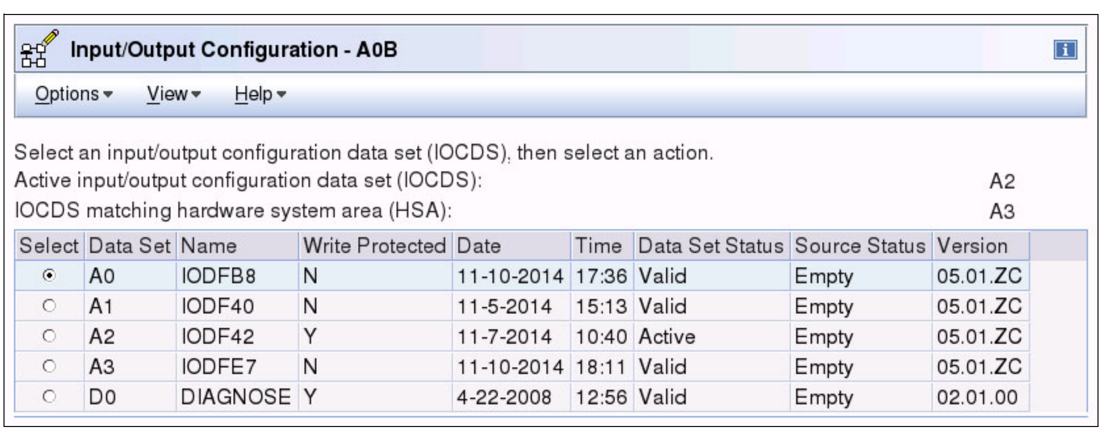

Figure 7-82 HMC - SAIOCP: IOCDS selection for import

11. Insert the USB flash memory drive that contains the IOCP text file. Wait for the drive insertion message to open (Figure 7-83).

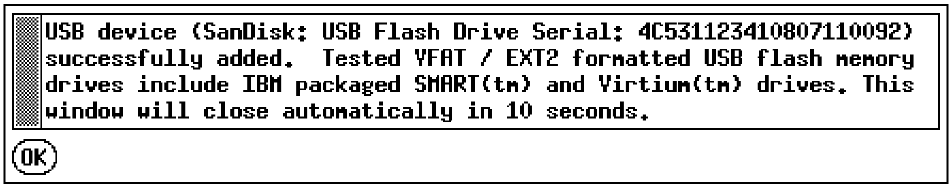

Figure 7-83 HMC - SAIOCP: USB drive insertion message

Tip: Only files in the root directory of the USB drive can be read by the HMC. Any folders and their contents will not be read.

12. Click **Options** → **Import Source File** → **Hardware Management Console USB Flash Memory Drive** (Figure 7-84).

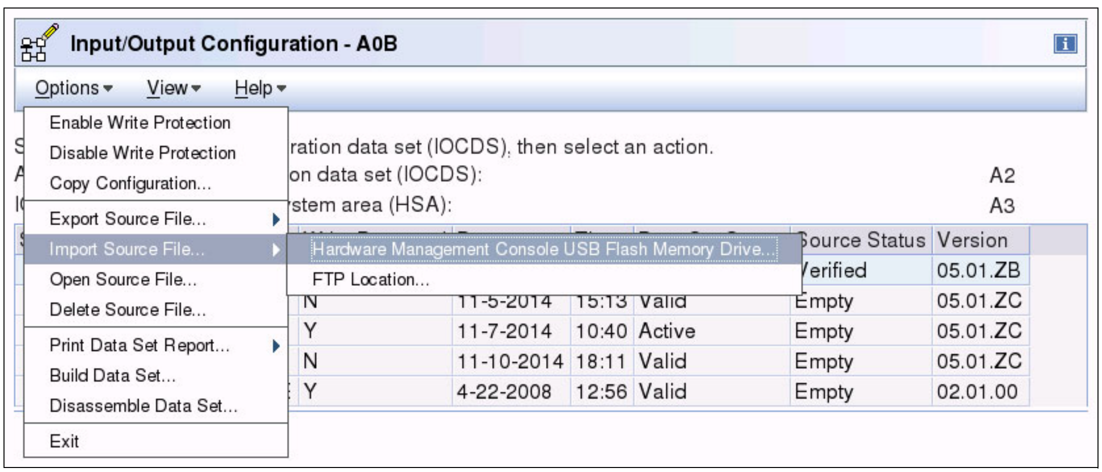

Figure 7-84 HMC - SAIOCP: import source file

13. Select the source file name and click **OK** (Figure 7-85).

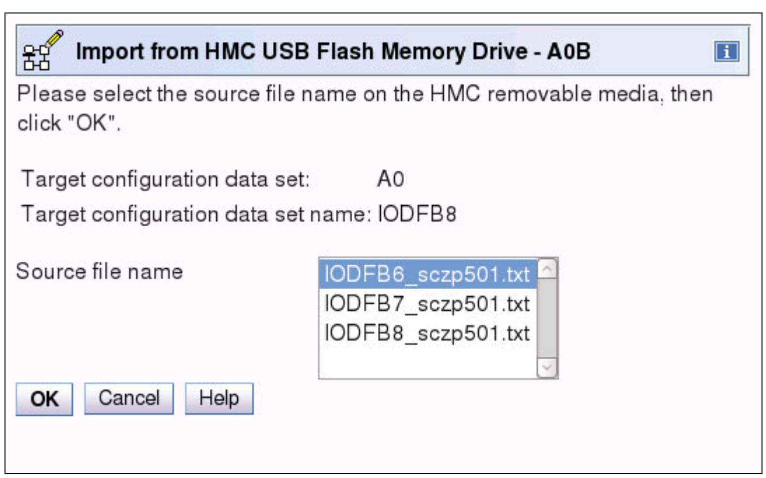

Figure 7-85 HMC - SAIOCP: select source file

14. The source file is now read from the USB drive. Click **OK** (Figure 7-86).

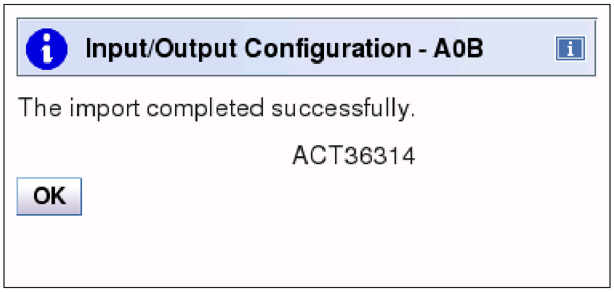

Figure 7-86 HMC - SAIOCP: source file imported

The Source Status now indicates `Imported` (Figure 7-87).

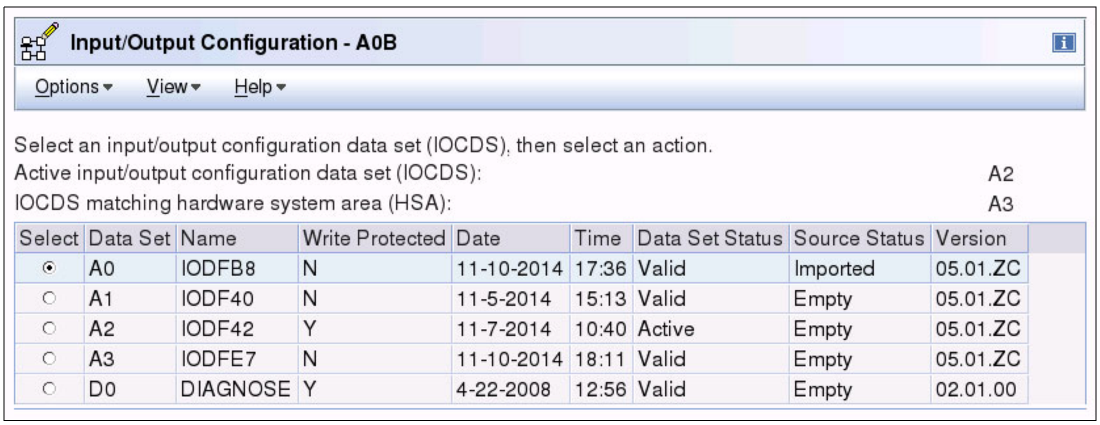

Figure 7-87 HMC - SAIOCP: IOCDS source status has changed to imported

15. Click **Options** → **Build Data Set** (Figure 7-88).

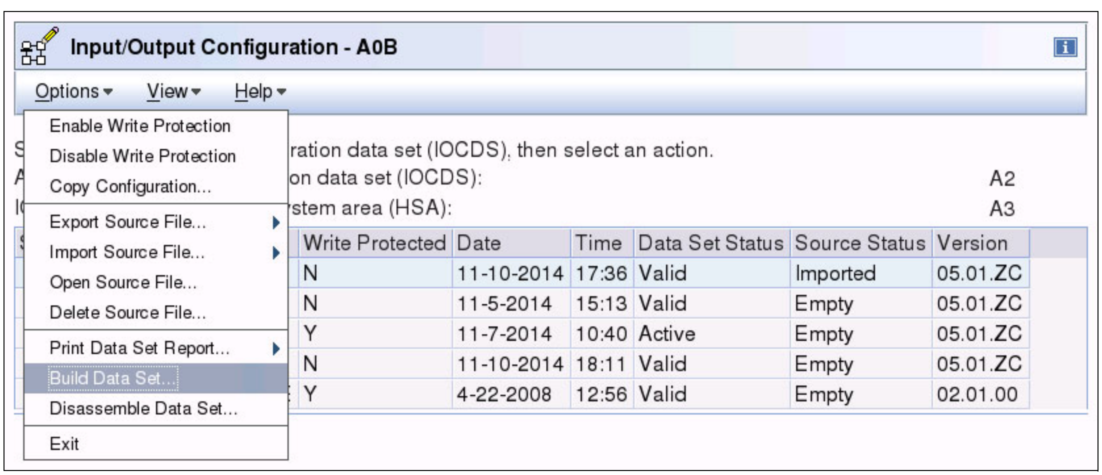

Figure 7-88 HMC - SAIOCP: building the IOCDS

16. Select the build options you want and click **OK** (Figure 7-89).

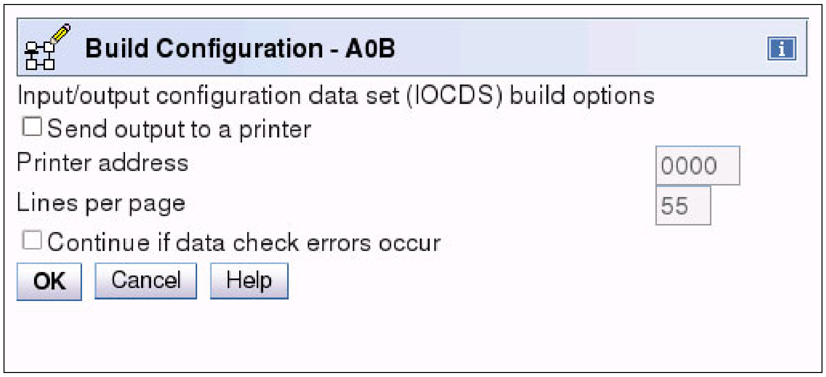

Figure 7-89 HMC - SAIOCP: build options

17. Observe the Build warning message, enter your HMC password, and click **Yes** (Figure 7-90).

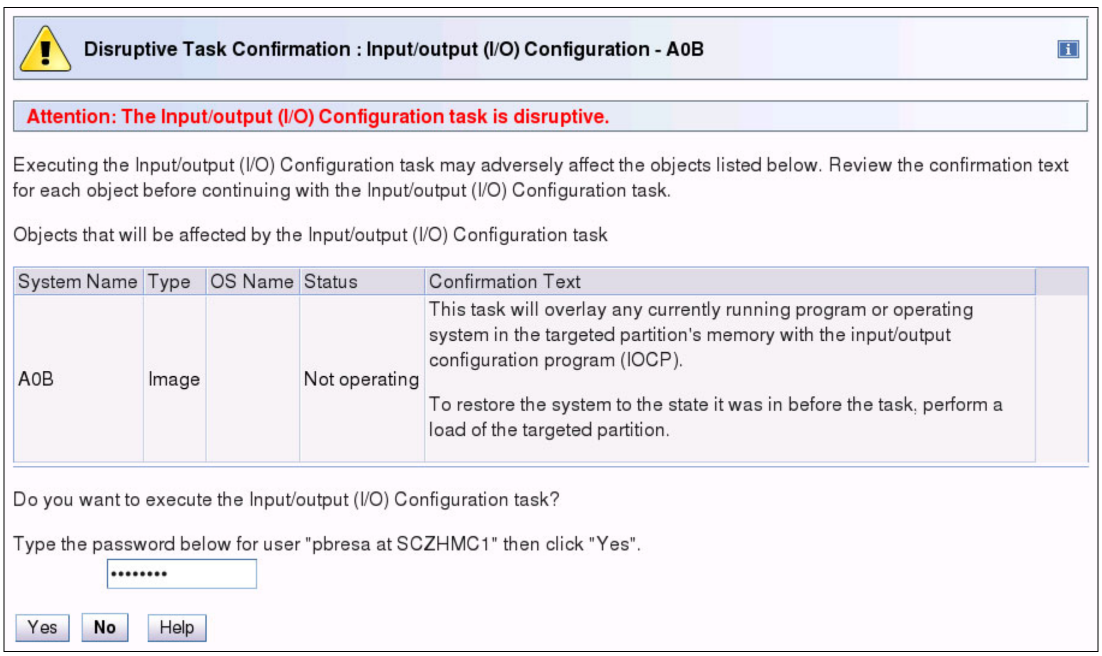

Figure 7-90 HMC - SAIOCP - build warning window

18. Status messages are displayed during the build process. After the process completes successfully, a message is displayed (Figure 7-91). Click **OK**.

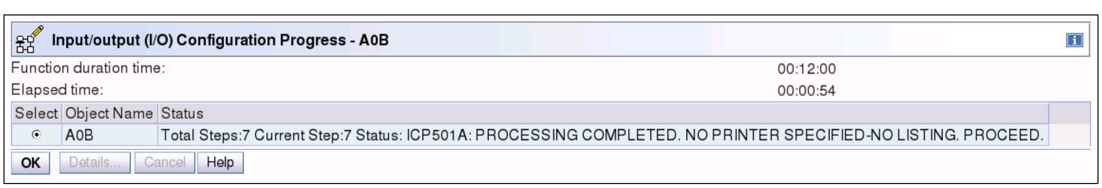

Figure 7-91 HMC - SAIOCP - build process status message

19.Observe that the Source Status now indicates `Verified` (Figure 7-92).

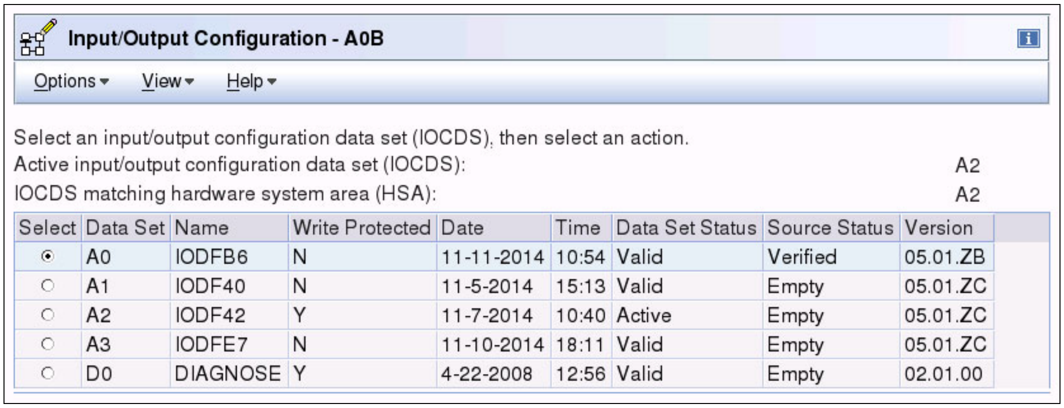

Select	Data Set	Name	Write Protected	Date	Time	Data Set Status	Source Status	Version
⦿	A0	IODFB6	N	11-11-2014	10:54	Valid	Verified	05.01.ZB
○	A1	IODF40	N	11-5-2014	15:13	Valid	Empty	05.01.ZC
○	A2	IODF42	Y	11-7-2014	10:40	Active	Empty	05.01.ZC
○	A3	IODFE7	N	11-10-2014	18:11	Valid	Empty	05.01.ZC
○	D0	DIAGNOSE	Y	4-22-2008	12:56	Valid	Empty	02.01.00

Figure 7-92 HMC - SAIOCP - IOCDS source status changed to verified

20.Click **Options** → **Exit** to end the IOCDS build process and deactivate the LPAR if it is no longer required.

This IOCDS is now ready to be selected by a Reset Profile. The 2964 can be activated (power-on reset) with the production IODF.

The USB drive can also now be removed. A drive removal message is displayed (Figure 7-93).

Figure 7-93 HMC - SAIOCP: USB drive removal message

7.9 HMC: Steps for profile activation

To activate your Reset profile using the HMC, complete the steps in this section.

7.9.1 Building the Reset Profile for activation 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 2964 after it is upgraded and handed over from the IBM System Service Representative.

To build the profile, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, or use a remote web browser and select the new 2964.
2. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
3. Under Systems or Members, click the system to select it (in this example, SCZP501).
4. On the Tasks tab, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (Figure 7-94 on page 416).

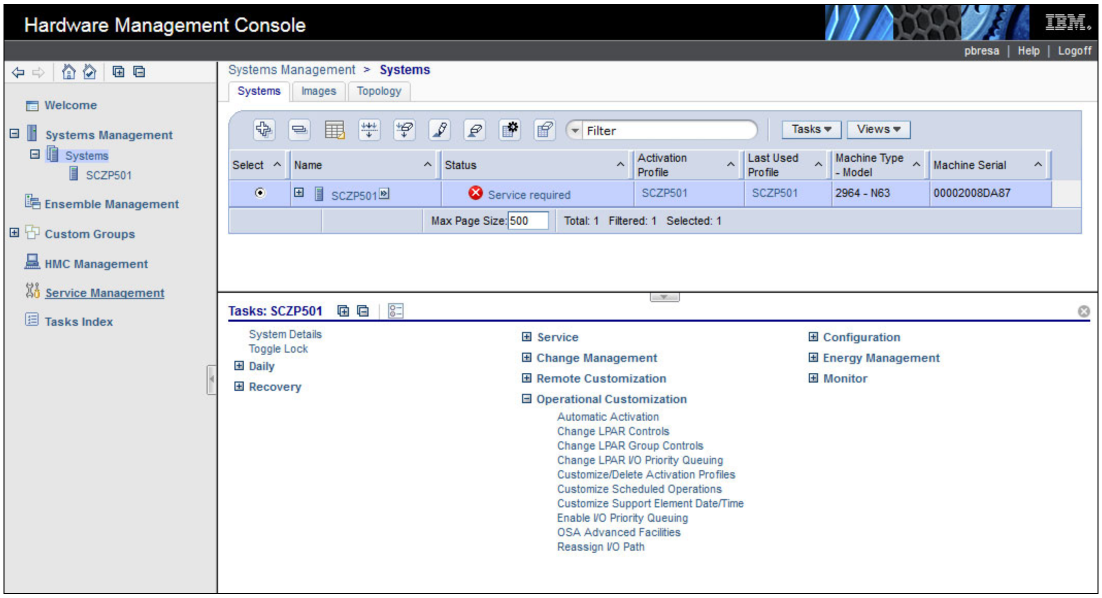

Figure 7-94 HMC - Customize Activation Profiles

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 this new TESTRESET profile and click **Customize profile**.
8. Click the IOCDS that you updated in 7.8, “HCD/HMC: Loading the 2964 processor IOCDS” on page 405. The ACTB0PDL message is displayed (Figure 7-95).

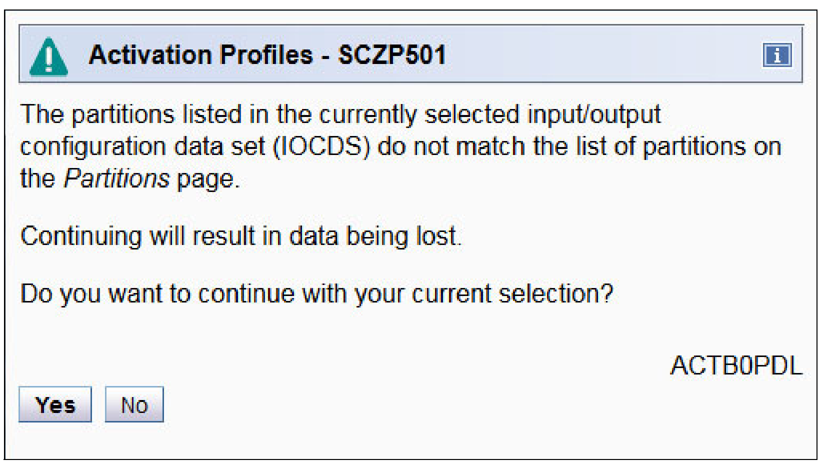

Figure 7-95 HMC - Activation Profiles: ACTB0PDL message

9. Depending on the circumstances, you can answer **Yes** or **No**. You might want to review the Partition Activation List now. For this example, we click **Yes.**
10. The HMC retrieves any Image profiles that match the LPAR names in the IOCDS that was selected. It also allows you to create new Image profiles for those that it cannot retrieve. In the example, the last option is selected. Click **OK** (Figure 7-96 on page 417).

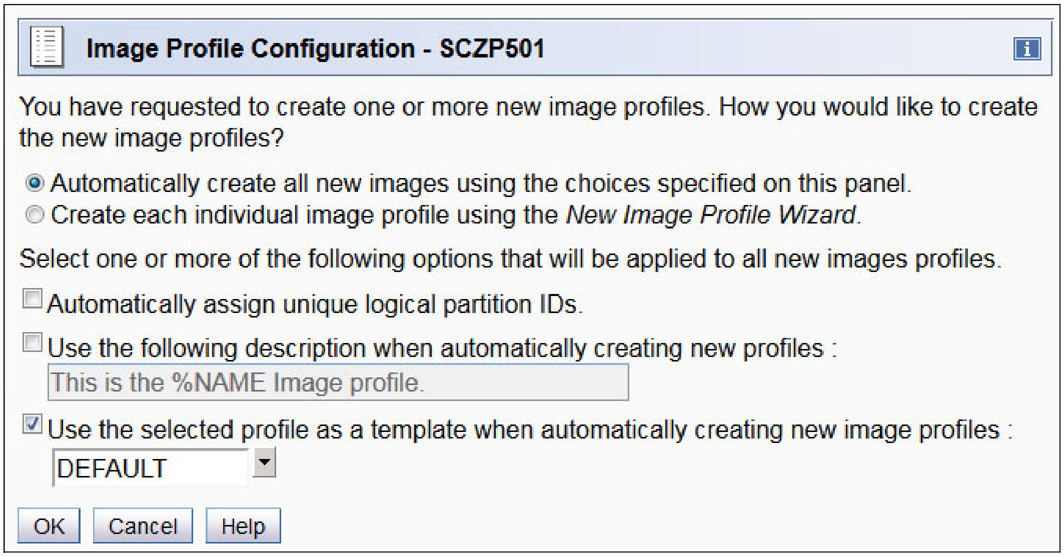

Figure 7-96 HMC - Image Profile automatic build options

11. Note the list of LPARs that were retrieved and built based on the LPARs that were defined in the selected IOCDS. Click **Save** (Figure 7-97).

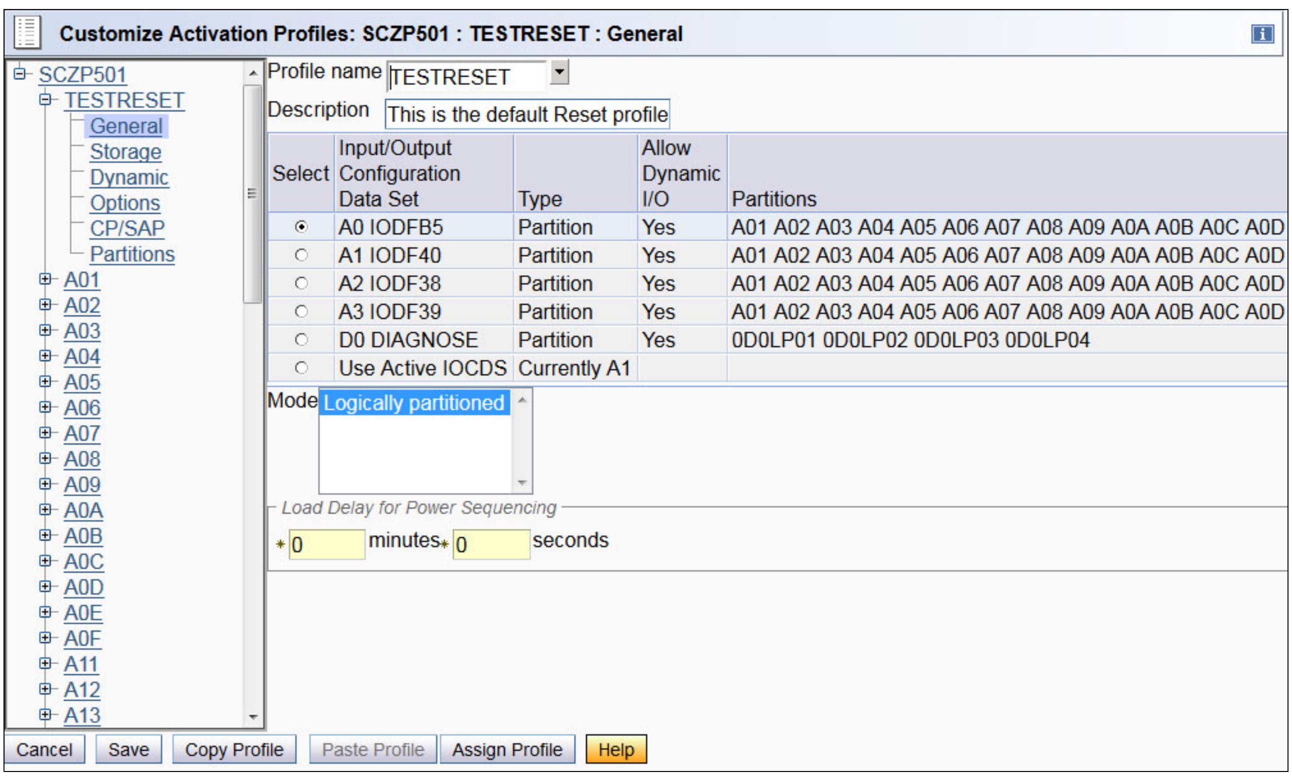

Figure 7-97 HMC - Reset and Image profile list

12. Review the items mentioned in these sections:

   – 7.9.2, “Building and verifying Image Profiles” on page 418
   – 7.9.5, “Server Time Protocol (STP) configuration” on page 418

7.9.2 Building and verifying Image Profiles

While still in the Reset Profile, you can review and update the Image Profile attributes.

7.9.3 Building and verifying Load Profiles

Create Load (IPL) Profiles by using the DEFAULTLOAD Load profile as a template for all the logical partitions for which you are performing an IPL on this processor.

7.9.4 Building and verifying Loadxx Members in SYS#.IPLPARM

You might need more Loadxx members defined in `SYS#.IPLPARM` for this processor.

If you used the HWNAME parameter to point to the Processor ID, update this parameter to point to the new Processor ID (in this example, from SCZP401 to SCZP501). Sometimes the LPARNAME parameter is also used in the Loadxx members and might need to be reviewed or updated.

7.9.5 Server Time Protocol (STP) configuration

Review the Server Time Protocol (STP) configuration to ensure that the correct External Time Source (ETS) was configured, and that you are connected to the correct Coordinated Time Network (CTN).

For more information about setting up your STP environment, see Chapter 9, “Server Time Protocol (STP)” on page 511.

7.9.6 Running a power-on reset (POR) of the 2964

When the 2817 or 2827 processor is upgraded with a 2964, the IBM System Service Representative will perform 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 in 7.9.1, “Building the Reset Profile for activation and pointing to required IOCDS” on page 415.

The 2964 is now ready to be Activated (power-on reset) using the Production Reset Profile.

8

Installing an additional IBM z13

This chapter describes how to add an IBM z13 into an existing IBM z Systems environment.

Because a wide variety of environments exists, your environment might not contain the same elements as the configuration described here. Nevertheless, the steps in this process provide enough information for you to replicate the approach in your own environment.

This chapter also shows how to define a coupling link connection between zEC12 and z13.

This chapter includes the following topics:

- Scenario overview
- HCD: Creating a 2964 work IODF
- HCD: Validating the 2964 work IODF
- CMT: Assigning PCHIDs to CHPIDs
- HCD: Updating the 2964 work IODF with PCHIDs
- HCD: Building the 2964 production IODF
- HCD/HMC: Loading the 2964 processor IOCDS
- HMC: Steps for profile activation

8.1 Scenario overview

This section describes the scenario to add an IBM z13 system.

8.1.1 The configuration process

The ten I/O configuration steps that are described in the "I/O configuration process" topic of *I/O Configuration Using z/OS HCD and HCM,* SG24-7804 are used for the example scenario.

Figure 8-1 shows the general process flow that is followed in this example. The numbered steps are described after the figure.

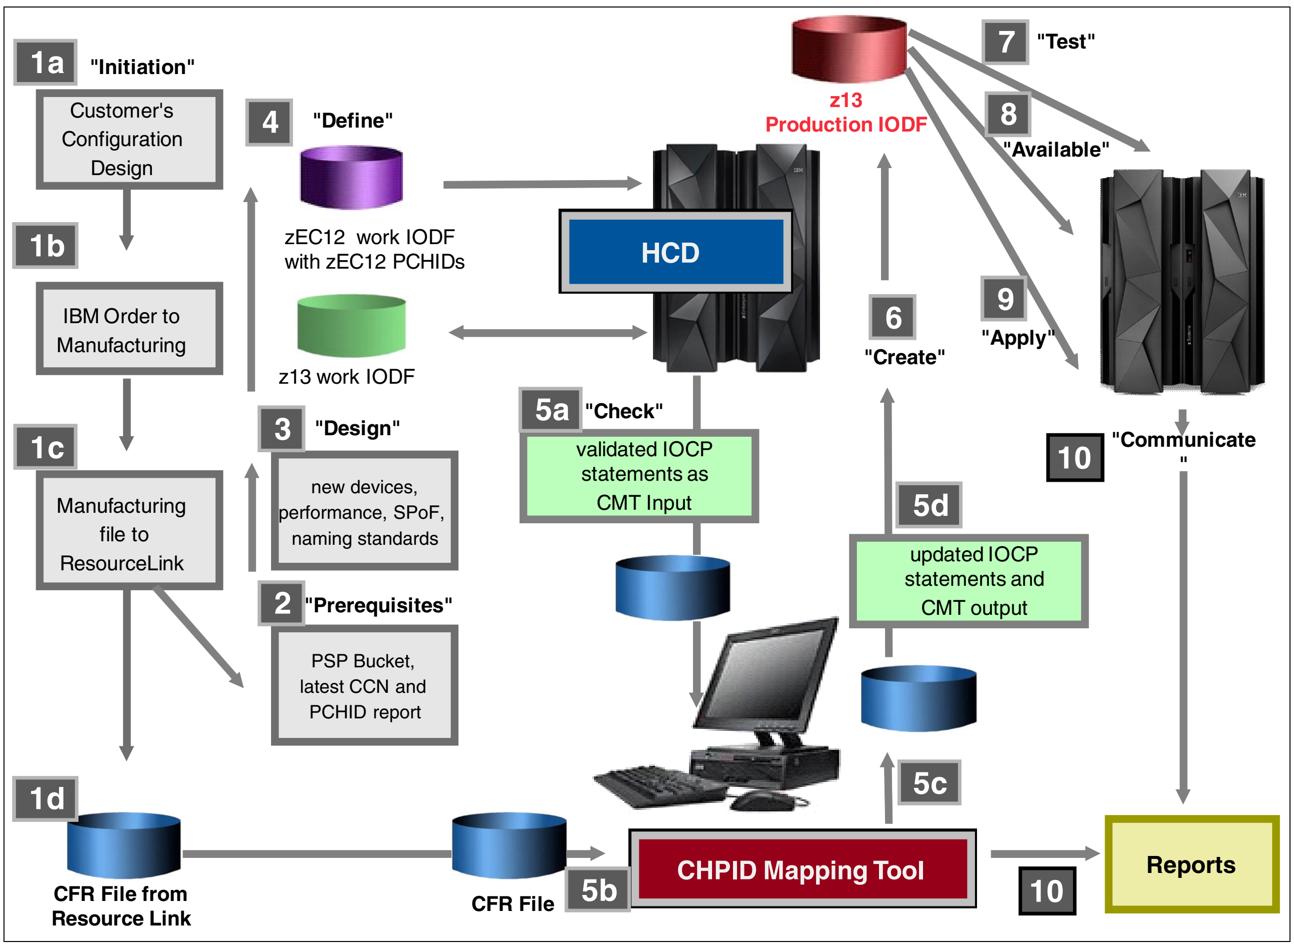

Figure 8-1 Overall configuration process flow

The steps in the figure are as follows:

1. Initiation

   a. When planning to migrate to an IBM z13, the IBM Technical Support team can help you define a configuration design that meets your needs. The configuration is then used during the ordering process.

   b. The IBM order for the configuration is created and passed to the manufacturing process.

c. The manufacturing process creates a configuration file that is stored at the IBM Resource Link website. This configuration file describes the hardware being ordered. This data is available for download by the client installation team.

d. A New Order report is created that shows the configuration summary of what is being ordered along with the Customer Control Number (CCN). The CNN can be used to retrieve `CFReport` (a data file that contains a listing of hardware configuration and changes for a central processor complex (CPC)) from Resource Link.

2. Prerequisites

   Ensure that you have the current PSP Bucket installed. Also, run the SMP/E report with fix category (FIXCAT) exceptions to determine whether any program temporary fixes (PTFs) must be applied. Ensure that you have the most current physical channel ID (PCHID) report and CCN from your IBM System Service Representative.

3. Design

   When you plan your configuration, consider this information:

   – Naming standards
   – FICON switch and port redundancy
   – Adequate I/O paths to your devices for performance
   – OSA channel-path identifier (CHPID) configuration for network and console communications
   – Coupling facility connections internally and to other systems.

   Because the IBM z13 server does not support attachment to the IBM Sysplex Timer, you must consider how the z13 will receive its time source. A z13 cannot join a Coordinated Time Network (CTN) that includes a z10 or earlier as a member. Because the z10 was the last server that supported the IBM Sysplex Timer (9037) connectivity, the z13 cannot be configured as a member of a mixed CTN. The z13 can only join an STP-only CTN.

   When you are planning to replace a z196 or zEC12 with a new z13, plan the replacement of channels that are not supported on IBM z13. You must carefully plan how to replace those, for instance, ISC-3 to HCA3-O or ICA SR for connectivity between z13 and another z13.

   You might need to increase CF storage size when you replace z196 or zEC12 with z13. Coupling Facility Control Code (CFCC) level 20 requirements might differ from CFCC level 19 and earlier. Use the CFSizer Tool to get the new CF storage requirements.

4. Define

   The existing z196 or zEC12 I/O configuration is used as a starting point for using Hardware Configuration Definition (HCD). The z196 or zEC12 production input/output definition file (IODF) is used as input to HCD to create a work IODF that becomes the base of the new IBM z13 configuration.

   When the new IBM z13 configuration is added and the obsolete hardware is deleted, a validated version of the configuration is saved in an IBM z13 validated work IODF.

5. Check

   a. From the validated work IODF, create a file that contains the IBM z13 IOCP statements. This IOCP statements file is transferred to the workstation used for the CHPID Mapping Tool (CMT). Hardware Configuration Manager (HCM) can also be used here to transfer the IOCP deck to and from the CMT.

   b. The configuration file that is created by the IBM Manufacturing process in step 1d is downloaded from Resource Link to the CMT workstation.

      The CHPID Mapping Tool (CMT) uses the input data from the files to map logical channels to physical ones on the new IBM z13 hardware.

You might have to make decisions in response to the following situations, among others:

i. Resolving situations in which the limitations on the purchased hardware cause a single point of failure (SPoF). You might must purchase more hardware to resolve some SPoF situations.

ii. Prioritizing certain hardware items over others.

c. After the CMT processing finishes, the IOCP statements contain the physical channels to logical channels assignment that is based on the actual purchased hardware configuration.

The CHPID Mapping Tool (CMT) also creates configuration reports to be used by the IBM System Service Representative and the installation team.

The file that contains the updated IOCP statements created by the CMT, which now contains the physical channels assignment, is transferred to the host system.

d. Use HCD, the validated work IODF file created in step 5a, and the IOCP statements updated by the CMT to apply the physical channel assignments created by the CMT to the configuration data in the work IODF.

6. Create

After the physical channel data is migrated into the work IODF, an IBM z13 production IODF is created and the final IOCP statements can be generated. The installation team uses the configuration data from the IBM z13 production IODF when the final power-on reset is done, yielding a z13 with an I/O configuration ready to be used.

7. Test

IODFs that are modifying existing configurations can be tested in most cases to verify that the IODF is making the intended changes.

8. Available

a. If you are upgrading an existing z196 or zEC12, you might be able to use HCD to write an IOCDS to your system in preparation for the upgrade. If you can write an IOCDS to your current system in preparation for upgrade, do so and let the IBM System Service Representative know which IOCDS to use.

> **Tip:** Using the HCD option "Write IOCDS" in preparation of an upgrade is the preferred method for writing the initial IOCDS when upgrading from a z196 or a zEC12 to a z13. This scenario uses the HCD option *Write IOCDS process*.

b. If the z196 or zEC12 is not network connected to the CPC where HCD is running, or if you are not upgrading or cannot write an IOCDS in preparation for the upgrade, use HCD to produce an IOCP input file. Download this input file to a USB flash drive.

9. Apply

The new production IODF can be applied to the IBM z13 in these ways:

– Using the power-on reset process
– Using the Dynamic IODF Activate process

10. Communicate

Communicating new and changed configurations to operations and the appropriate users and departments is important.

8.1.2 Planning considerations

The following I/O features can be ordered for a new IBM z13:

- FICON Express16S LX (long wavelength - 10km) - FC 0418
- FICON Express16S SX (short wavelength) - FC 0419
- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Crypto Express5S - FC 0890
- Flash Express - FC 0403
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171
- PCIe-O SR Integrated Coupling Adapter - FC 0172

The following features, if present in a z196 or zEC12, can be carried forward when you upgrade to an IBM z13:

- FICON Express8S LX (long wavelength - 10 km) - FC 0409
- FICON Express8S SX (short wavelength) - FC 0410
- FICON Express8 LX (long wavelength - 10 km) - FC 3325
- FICON Express8 SX (short wavelength) - FC 3326
- OSA-Express5S GbE LX (long wavelength) - FC 0413
- OSA-Express5S GbE SX (short wavelength) - FC 0414
- OSA-Express5S 10 GbE LR (long reach) - FC 0415
- OSA-Express5S 10 GbE SR (short reach) - FC 0416
- OSA-Express5S 1000BASE-T Ethernet - FC 0417
- OSA-Express4S GbE LX (long wavelength) - FC 0404
- OSA-Express4S GbE SX (short wavelength) - FC 0405
- OSA-Express4S 10 GbE LR (long reach) - FC 0406
- OSA-Express4S 10 GbE SR (short reach) - FC 0407
- OSA-Express4S 1000BASE-T Ethernet - FC 0408
- RoCE Express 10 GbE SR (short range) - FC 0411
- zEDC Express - FC 0420
- Flash Express - FC 0402
- HCA3-O 1x LR IFB - FC 0170
- HCA3-O 12x IFB - FC 0171

The following features are not supported on IBM z13:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet

- OSA-Express2 10 GbE Long Reach
- OSA-Express
- Crypto Express3
- Crypto Express2
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode
- ICB-2
- ICB-3
- ICB-4
- PCIXCC and PCICA
- Parallel channels (use FICON and ESCON converters)

Table 8-1 lists the channel types as described in an input/output configuration data set (IOCDS) that are supported by the z196 (2817), zEC12 (2827), and z13 (2964).

Table 8-1 Channels, links, and adapters with CHPID type and support

Channels	CHPID type	2817 support	2827 support	2964 support
ESCON channels: Connection Channel (ESCON architecture) Channel to Channel (connects to CNC) Converter Channel Path (for BL types) Converter Channel Path (for BY types)	CNC CTC CVC CBY	Up to 240 Up to 240 Up to 240 Up to 240	- No - -	- No - -
FICON native channels that attach to FICON directors or directly to FICON control units: FICON Express 4 SX and LX FICON Express 8 SX and LX FICON Express 8S SX and LX FICON Express 16S SX and LX	FC FC FC FC	Carry forward up to 288 Yes No	Carry forward Carry forward up to 320 No	No up to 64 Carry Forward up to 320
FICON channels that attach to Fibre Channel devices, switches, directors, or Fibre-Channel-to-SCSI bridges	FCP	Yes	Yes	Yes
ISC-3 peer mode channels (connects to another CFP)	CFP	Up to 48	Carry forward	No
IC peer channels (connects to another ICP)	ICP	Up to 32	Up to 32	Up to 32
HCA2-O 12x InfiniBand host channel adapters **HCA2-O 1x LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 32	Up to 32 links Up to 32 links	No No
HCA3-O 12x InfiniBand host channel adapters **HCA3-O LR** InfiniBand host channel adapters	CIB	Up to 32 Up to 48	Up to 32 links Up to 64 links	Up to 32 links Up to 64 links
PCIe-O SR Integrated Coupling Adapter	CS5	No	No	Up to 32 links
HiperSockets (IQDIO) channels	IQD	Up to 32	Up to 32	Up to 32
OSA- Express2 GbE LX/SX	OSD and OSN	Up to 48 ports carried forward	No	No
OSA- Express2 1000BASE-T	OSE, OSD, OSC, and OSN	Yes	No	No
OSA- Express2 10 GbE LR	OSD	Yes	No	No

Channels	CHPID type	2817 support	2827 support	2964 support
OSA-Express3 GbE LX/SX	OSD and OSN	Up to 96 ports	Carry forward	No
OSA-Express3 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	OSC, OSD, OSE, and OSN: Yes OSM: No	Carry forward	No
OSA- Express3 10 GbE LR/SR	OSD and OSX	OSD: Yes OSX: No	Carry forward	No
OSA-Express4S GbE LX/SX	OSD	Up to 96 ports	Up to 96 ports	Carry forward
OSA-Express4S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Carry forward
OSA-Express4S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Carry forward
OSA-Express5S GbE LX/SX	OSD	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 1000BASE-T	OSE, OSD, OSC, OSN, and OSM	No	Up to 96 ports	Up to 96 ports
OSA-Express5S 10 GbE LR/SR	OSD and OSX	No	Up to 48 ports	Up to 48 ports
10GbE RoCE Express SR	ROCE	No	Up to 32 ports	Up to 32 ports
zEDC Express Enterprise Data Compression	ZEDC	No	Max 8 features	Max 8 features

Consider the following information when you plan your configuration:

- Coupling links
- Hardware Management Console (HMC)
- Software support
- Open Systems Adapter - Integrated Console Controller
- Fibre Channel Protocol
- CPC name versus Processor ID
- Local system name

Coupling links

Only the following coupling facility CHPIDs are supported:

- CHPID type CIB: PSIFB links connecting to an HCA3-O 12x and HCA3-O LR 1x
- CHPID type ICP: Internal Coupling links.
- CHPID type CS5: Integrated Coupling adapter PCIe-O SR

Considerations: Coupling links can be defined as both Coupling and STP links, or STP-only links. IBM z13 does not support the ISC-3 features.

Hardware Management Console (HMC)

The HMC can appear either as the current HMC does, or as an HMC that can run code to manage an Ensemble. The current HMC is used to manage, monitor, and operate one or more IBM z Systems servers and their associated logical partitions. An HMC that has ensemble code running is an HMC attached to one or more zEnterprise Systems configured as ensemble members. A particular ensemble is managed by a pair of HMCs in primary and alternate roles.

The HMC has a global (Ensemble) management function, whereas the SE has local node management responsibility. When tasks are performed on the HMC, the commands are sent to one or more Support Elements (SEs), which then issue commands to their CPCs.

The IBM z13 requires HMC Application V2.13.0 (driver level 21) or later, and uses only Ethernet for its network connection. The HMC and the SEs do not contain a floppy disk drive. Therefore, you must use a USB flash memory drive to input and back up client configuration data.

Software support

HCD V2.1 and later with the Preventive Service Planning (PSP) bucket for 2964DEVICE and PTFs) is required to define and support some of the new features of the IBM z13.

Open Systems Adapter - Integrated Console Controller

Because support has been withdrawn for the 2074 console controllers, consider using OSA-Express4S 1000BASE-T or OSA-Express5S1000BASE-T CHPIDs defined as TYPE=OSC. With this OSA card, you can set up console function that is supported by a configuration file defined on the Support Element for that processor.

Fibre Channel Protocol

When you use CHPIDs defined as TYPE=FCP, consider N-Port ID Virtualization (NPIV).

For more information about FCP CHPIDs and the WWPN prediction tool to manage them, see “Batch Network Analyzer and Compression” on page 47.

CPC name versus Processor ID

HCD allows you to define different processors (logical) to the same CPC (physical). The Processor ID must be unique within the same IODF, but the CPC name does not. Therefore, the CPC name does not need to match the Processor ID. This advantage is useful when you have several processor/logical partition/control unit setups that share a physical CPC within the same IODF. Furthermore, the Processor ID is what is defined for the optional HWNAME parameter in the LOAD member in SYS1.IPLPARM.

The CPC name is coded in HCD option 1.3 under View Processor Definition in the CPC name field under SNA address, along with a Network name. It is the CPC name, and not the Processor ID, that is displayed on the HMC.

When you view the network information for a CPC through the HMC, the SNA address is made up of a network name and CPC name separated by a dot (for example, `USIBMSC.SCZP501`). These values are defined in the Support Element for the CPC. They must match the values that are set in the IODF so that HCD option 2.11 can find the CPC to write an IOCDS in the System z Cluster List.

Local system name

An extra system name, LSYSTEM, is used to identify the local system name of a server when you define PSIFB type=CIB or ICA type=CS5 coupling links.

This data field can be found when you change a CIB-capable processor under HCD option 1.3.

The LSYSTEM field can be set or changed to any one to eight alphanumeric characters. Also, it can begin with either an alphabetic or numeric character. All characters are uppercase.

The following rules determine whether, and where, HCD sets the LSYSTEM keyword automatically:

- If a CIB-capable processor is defined and the CPC name is set but the local system name is not set, HCD sets the local system name to the CPC name.
- If a CIB-capable processor that has not yet defined a CPC name is changed to obtain a CPC name but no local system name, HCD sets the CPC name to the local system name.
- If a non-CIB capable processor is changed to a support level that is CIB capable, and the processor has a CPC name set but no local system name, the local system name defaults to the CPC name.
- If the processor definition is changed such that the local system name is explicitly removed, HCD does not do any defaulting.
- If a processor has a local system name set (whether it has a CPC name or not), any change to the local system name must be done explicitly. There is no implicit name if the CPC name or the support level is changed.
- During Build Production IODF, verification determines whether a local system name is set for the processor if the processor has a CIB channel path defined. If this verification fails, an error message is displayed and the production IODF is not built.

Generally, set the local system name the same as the CPC name.

The following are more keywords for the ID statement in an IOCP deck:

CPATH Specifies the CSSID and CHPID on the connecting system.

CSYSTEM Connecting System name (LSYSTEM in processor definition).

AID Host Card Adapter ID (found in PCHID report).

PORT Host Card Adapter port.

8.1.3 Installing an extra 2964

This scenario shows configuration steps for installing a new 2964 processor into an existing hardware environment. This process has the following key considerations:

- HCD requires a new processor ID for the 2964.
- HCD requires a new CPC name for the 2964.
- The 2964 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 production IODF.
- The target IODF is a new 2964 work IODF.

This example shows a new 2964-N63 with a Processor ID of SCZP501 and with six CSSs (CSS ID=0, CSS ID=1, CSS ID=2, CSS ID =3, CSS ID=4, and CSS ID=5). The CPC name of SCZP501 and serial number of 02-8DA87 are used for the 2964.

The following CHPID types are migrated:

- OSD, OSC, OSM and OSX
- FC and FCP
- CIB and ICP
- IQD

The following hardware/CHPID types are *not* supported and *not* migrated to the 2964:

- ESCON
- FICON Express4 LX (long wavelength - 10 km)
- FICON Express4 SX (short wavelength)
- FICON Express2 (LX and SX)
- FICON Express (LX and SX)
- FICON (pre-FICON Express)
- OSA-Express3 GbE LX (long wavelength)
- OSA-Express3 GbE SX (short wavelength)
- OSA-Express3 10 GbE LR (long reach)
- OSA-Express3 10 GbE SR (short reach)
- OSA-Express3 1000BASE-T Ethernet
- Crypto Express3
- HCA2-O 1x LR IFB
- HCA2-O 12x IFB
- ISC-3 (Peer mode only)
- ISC-3 Links in Compatibility Mode

Table 8-2 summarizes the tool requirements. The process is documented later in this chapter.

Table 8-2 Extra 2964 I/O configuration

Extra 2964	**Extra 2964s to connect to the same switch ports and control units to which existing processors connect (new 2964s)**
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 (see 8.2, “HCD: Creating a 2964 work IODF” on page 429)
CMT Program	Optional, but generally a good idea
`CFReport File` (CCN)	Required
IOCP (import from validated work IODF)	Yes
CMT Actions (PCHID reset)	Yes
CMT IOCP Output	Yes
CMT Reports	Yes, CHIPID Report and CHIPID to CU Report

8.2 HCD: Creating a 2964 work IODF

The following steps explain how to define a 2964 processor to the existing I/O configuration by using HCD:

1. Create a 2964 work IODF from a current production IODF.
2. Add the processor.
3. Change the required partition names and usage from reserved.
4. Add CHPIDs.
5. Connect FICON CHPIDs to existing switches and available ports.
6. Create any additional control units unique to this processor, and connect CHPIDs.
7. Define all required CF connections to other processors and any internal CF connections required.
8. Create FICON channel-to-channel connections (FCTCs).
9. Connect to existing control units for DASD, Tape, Printer, and Communication devices directly or over switches.
10. Create the OSA Configuration (OSC, OSD, OSE, OSN, OSX, and OSM).
11. Build a validated work IODF.
12. Create an IOCP statements file and transfer it to your CHPID Mapping Tool workstation. This step can also be completed with HCM.
13. Import CFReport and IOCP statements into the CMT.
14. Perform hardware resolution and PCHID/CHIPID availability.
15. Create configuration reports for yourself and the IBM System Service Representative.
16. Import IOCP statements that are updated with PCHIDs back into the validated work IODF.
17. Build a production IODF.
18. Remotely write the IOCP to an IOCDS on the 2964 Support Element. If this action is not possible, copy the IOCP statements to a USB memory flash drive and run the Stand-Alone Input/Output Configuration Program to load the IOCP statements to an IOCDS on the 2964 Support Element.
19. Build Reset, Image, and Load profiles if required.
20. Run a power-on reset (activate) of the 2964.

The following sections describe some of these steps in more detail. For more information about defining processors, partitions, switches, control units, and devices in an IODF by using HCD, see *z/OS HCD Users Guide*, SC33-7988 and *IOCP Users Guide*, SB10-7037.

8.2.1 Creating a work IODF from the current production IODF

In HCD, select the current production IODF that contains the existing hardware environment that will be connected to the new 2964 (for example, `SYS6.IODFB2`).

8.2.2 Adding the new 2964 processor

To add the new 2964 processor, complete the following steps:

1. From the HCD main menu, select option **1.3, Processor List**. I
2. n the Processor List (Figure 8-2), 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
******************************* Bottom of data ********************************
```

Figure 8-2 HCD - Processor List: adding a processor

The Add Processor panel opens (Figure 8-3).

```
------------------------------ 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  . . . . . . . ________
```

Figure 8-3 HCD - Add Processor: data fields to be updated

3. Specify the appropriate values. For example, specify these values, shown in Figure 8-4 on page 431:

Processor ID	SCZP501
Processor type	2964
Processor model	N63
Number of channel subsystems	(Keep this blank for now)
Serial number	08DA872964
Network name	USIBMSC
CPC name	SCZP501
Local System Name	(Keep this blank for now)

```
----------------------------- Add Processor --------------------------


Specify or revise the following values.

Processor ID . . . . . . . . . . SCZP501_
Processor type . . . . . . . . . 2964____  +
Processor model  . . . . . . . . N63_____  +
Configuration mode . . . . . . . LPAR      +
Number of channel subsystems . . _         +

Serial number  . . . . . . . . . 08DA872964
Description  . . . . . . . . . . Sphinx______________________

Specify SNA address only if part of a System z cluster:

Network name . . . . . . . . . . USIBMSC_  +
CPC name . . . . . . . . . . . . SCZP501_  +

Local system name  . . . . . . . ________
```

Figure 8-4 HCD - 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.IODFB8.WORK`).
5. Press Enter. You now have an extra 2964 processor named SCZP501 (Figure 8-5).

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

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


 / Proc. ID Type +   Model +  Mode+ Serial-# + Description
 _ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
 _ SCZP401  2827     H43      LPAR  00B8D72827 Helix
 _ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
 ******************************* Bottom of data *******************************




+-----------------------------------------------------------------+
| Definition of processor SCZP501 has been extended to its maximum |
| configuration.                                                   |
+-----------------------------------------------------------------+
```

Figure 8-5 HCD - 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 is 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 2964, HCD automatically defines the maximum configuration of six CSSs and 84 logical partitions.

6. Enter `s` next to the SCZP501 and press Enter. The Channel Subsystem List is displayed. Here you can see six channel subsystems (CSS0-CSS5) that are defined with the default MAXDEV values for SS0 of 65280 set by HCD and 65535 set for SS1, SS2 and SS3 (Figure 8-6).

> **Tip:** The z13 and the zEC12 support IPL from subchannel set 1 (SS1), subchannel set 2 (SS2), or subchannel set 3 (SS3) in addition to subchannel set 0. 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 defined using the same device number and a new device type in an alternate subchannel set to be used for IPL, I/O definition file (IODF), and stand-alone dump volumes when needed.
>
> IPL from an alternate 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.

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

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

Processor ID . . . : SCZP501

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

Figure 8-6 HCD - Channel Subsystem List: four subchannel sets

7. Enter **p** next to one more CSS and observe that HCD also defines the maximum number of logical partitions, 15 per CSS (and 10 in CSS5) to make a total of 85, as Reserved (*), as shown in Figure 8-7.

```
---------------------------- Partition List ----------------------------
  Goto  Backup  Query  Help
-----------------------------------------------------------------------
                                                              Row 1 of 15
Command ===> ________________________________________ Scroll ===> CSR

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

Processor ID  . . . . : SCZP501
Configuration mode  . : LPAR
Channel Subsystem ID  : 0

/ Partition Name   Number Usage + Description
_ *                1      CF/OS   ________________________________
_ *                2      CF/OS   ________________________________
_ *                3      CF/OS   ________________________________
_ *                4      CF/OS   ________________________________
_ *                5      CF/OS   ________________________________
_ *                6      CF/OS   ________________________________
_ *                7      CF/OS   ________________________________
```

Figure 8-7 HCD - Partition List

Scroll to see the remaining reserved logical partitions that are defined for this CSS.

8. Define the resources to the new 2964 processor:
   a. Change Reserved partitions for each CSS to the required partition name and usage. Not all partitions must be changed now. They can be modified later by using the Dynamic IODF Activate process.
   b. Add CHPIDs to each CSS, with no PCHIDs assigned.
   c. Define a Partition Access List for these CHPIDs.
   d. Define a Partition Candidate List for these CHPIDs.

> **More information:** For more information about defining processors, partitions, switches, control units, and devices, see the *z/OS HCD Users Guide,* SC33-7988, and the *IOCP Users Guide,* SB10-7037.

9. Connect FICON CHPIDs to available ports on existing switches.
10. Connect FICON CHPIDs directly to control units, where applicable.
11. Create any control units unique for this processor and connect CHPIDs.
12. Define all required coupling links to other processors in the hardware environment, and also any internal coupling links required.
13. Create FICON CTCs.
14. Define logical paths to existing control units for DASD, tape, printers, and communications controllers directly or over the FICON switches.
15. Create OSA resources (such as OSC, OSD, OSE, OSN, OSM, and OSX).

8.2.3 Overdefining channel paths on an XMP processor

Sometimes you must define a channel path that is not physically installed on the processor. This definition is useful if you are planning to add more channel cards to the processor in the future and want to have the definitions in the IODF before the hardware is installed.

HCD allows you to overdefine CHPIDs by using an asterisk (*) for the PCHID value. An overdefined CHPID must adhere to all validation rules, but it is not taken into account by an IOCDS download. Also, it is not included in the IOCP statements, in a CONFIGxx member, or during dynamic activation.

If a control unit contains only CHPIDs with a PCHID value of an asterisk (*), the whole control unit (including any attached devices) is omitted from the configuration to be activated.

When you install the channel path later, you must edit the CHPID and replace the asterisk (*) with its valid PCHID.

Remember: This is not the case for CIB type CHPIDs, where these CHPIDs have connections to other CIB type CHPIDs. Therefore, HCD allows you to define CIB type CHPIDs as overdefined only if they are unconnected.

Overdefining is also supported for CS5 type CHPID definitions.

The 2964 production IODF can then be activated dynamically and the PCHID, CHPID, and control unit definitions become available to the operating system.

Figure 8-8 shows the CHPID/PCHID definitions *before* they are defined as overdefined. Press PF20 (right) in the Channel Path List.

```
                              Channel Path List     Row 1 of 107 More: <
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    1A3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    553   FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    153   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 8-8 HCD - Channel Path List: reserving CHPIDs

Figure 8-9 shows the CHPID/PCHID definitions *after* they are defined as overdefined.

```
                              Channel Path List    Row 22 of 109 More:      >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID : SCZP501    CSS ID  : 0
1=OS A01       2=OS A02       3=OS A03       4=OS A04       5=OS A05
6=OS A06       7=OS A07       8=OS A08       9=OS A09       A=OS A0A
B=OS A0B       C=CF A0C       D=CF A0D       E=CF A0E       F=CF A0F
        PCHID                   I/O Cluster  --------- Partitions 0x -----
/ CHPID AID/P Type+ Mode+ Mng   Name +       1 2 3 4 5 6 7 8 9 A B C D E F
_ 5F    271   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 60    1D1   FC    SPAN  No    ________     a a a a a a a a a a a _ _ _ _
_ 61    5C3   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 62    *     FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
_ 63    *     FC    SPAN  No    ________     a a a a a a a a a a _ _ _ _ _
_ 64    *     FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 65    2A3   FC    SPAN  No    ________     _ _ _ _ _ _ _ a a _ _ _ _ _ _
_ 66    243   FC    SPAN  No    ________     _ _ _ _ _ a _ _ _ a _ _ _ _ _
```

Figure 8-9 HCD - Channel Path List: overdefined CHPIDs

8.2.4 Adding CHPID type CIB channels and PSIFB links

This section addresses defining CIB CHPIDs between a 2817 or 2827 and the 2964, and defining a PSIFB link between the CIB type CHPIDs.

The local system name or LSYSTEM keyword is used for the PSIFB link definitions.

Requirement: The local system name must be defined if PSIFB links are going to be defined over CIB type CHPIDs. Generally, use the CPC name of that server.

To add the channels and links, complete the following steps:

1. From the HCD main menu, select option **1.3. Processor List** (Figure 8-10).
2. Enter **c** next to SCZP401 (the 2827) 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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
c SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
****************************** Bottom of data *******************************
```

Figure 8-10 HCD - Processor List: changing a processor definition

3. The local system name field is blank for this 2827. The data field is automatically set according to the following rules:
   - If a 2827 processor is defined and the CPC name is set, but the local system name is not set, HCD defaults the local system name to the CPC name.
   - If a 2827 processor that has not yet defined a CPC name is changed to obtain a CPC name but no local system name, HCD defaults the CPC name to the local system name.
   - If the 2827 processor definition is changed such that the local system name is explicitly removed, HCD does not do any defaulting.
   - If a 2827 processor has a local system name set (whether it has a CPC name or not), any change to the local system name must be done explicitly. There is no implicit name if the CPC name or the support level is changed.

   Generally, set the local system name so that it is the same as the CPC name. See Figure 8-11.

```
----------------------- Change Processor Definition ---------------


Specify or revise the following values.

Processor ID . . . . . . . . : SCZP401
Support level:
XMP, 2827 GA2 support
Processor type . . . . . . . . 2827      +
Processor model  . . . . . . . H43       +
Configuration mode . . . . . . LPAR      +

Serial number  . . . . . . . . 00B8D72827  +
Description  . . . . . . . . . Helix

Specify SNA address only if part of an System z cluster:

Network name . . . . . . . . . USIBMSC   +
CPC name . . . . . . . . . . . SCZP401   +

Local system name  . . . . . . ________
```

Figure 8-11 HCD - Change Processor Definition

4. In the example, add the local system name (Figure 8-12).

```
---------------------- Change Processor Definition -------------


Specify or revise the following values.

Processor ID . . . . . . . . : SCZP401
Support level:
XMP, 2827 GA2 support
Processor type . . . . . . . . 2827      +
Processor model  . . . . . . . H43       +
Configuration mode . . . . . . LPAR      +

Serial number  . . . . . . . . 00B8D72827  +
Description  . . . . . . . . . Helix

Specify SNA address only if part of an System z cluster:

Network name . . . . . . . . . USIBMSC   +
CPC name . . . . . . . . . . . SCZP401   +

Local system name  . . . . . . SCZP401_
```

Figure 8-12 HCD - Change Processor Definition: updating the local system name

Add CIB type CHPIDs to the 2827 and 2964 processors and define a PSIFB link by completing the following steps:

1. Review the PCHID report that came with the 2827 and look for the AID. The example uses AID=0A Port=2 (HCA 3-O).
2. Review the PCHID report that came with the 2964 and look for the AID. The example uses AID=08 Port=2 (HCA 3-O).

   The Adapter IDs for the Host Channel Adapter F/C 0171 on the 2964 are shown in Figure 8-13.

   Tip: After they are installed, Adapter IDs can also be found on the HMC/SE panels.

```
CHPIDSTART
  12345675                     PCHID REPORT                     Nov 10,2014
  Machine: 2964-N63 SNXXXXXXX
  - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
- -
  Source           Cage  Slot  F/C    PCHID/Ports or AID         Comment
  19/07/J01        A19A  LG07  0171   AID=08
  19/09/J01        A19A  LG09  0171   AID=0A
```

Figure 8-13 PCHID - 2964 PCHID Report: showing Adapter IDs

3. In the processor list, enter **s** next to the 2827 and press Enter.
4. Enter **s** next to the CSS where you want to define the CIB CHPIDs and press Enter.

5. Enter **add** command or press PF11 (add) to add new CHPIDs, and press Enter (Figure 8-14).

```
                              Channel Path List      Row 1 of 88 More:       >
Command ===> add                                             Scroll ===> CSR

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

Processor ID . . . . : SCZP401      Helix
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 00    574   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT All LPARs 9.12.4.x
_ 01    54C   OSC   SPAN  __  __ __       No  Exp4S 1KBaseT All LPARs 9.12.4.x
_ 02    58C   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT DSTPGZ/LABSERV
_ 03    5B8   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT COMMPLEX/VMLINUX9
_ 04    534   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT COMMPLEX/LABSERV
_ 06    50C   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT ECS 012x
_ 07    570   OSD   SPAN  __  __ __       No  Exp4S 1KBaseT ECS 014x
```

Figure 8-14 HCD - Channel Path List: adding CHPIDs

6. Enter the appropriate values. For example, specify the following values, as shown in Figure 8-15:

Channel path ID 90
Channel path type CIB
Operational mode SHR

```
------------------------------ Add Channel Path ------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP401      Helix
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 90    +           Channel ID ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CIB   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +
Description  . . . . . . ________________________________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 8-15 HCD - Add Channel Path: defining a CIB CHPID

7. Press Enter. HCD now displays the Specify HCA Attributes panel.
8. Enter the appropriate values. As shown in Figure 8-16, the example uses the following values:

   Adapter ID of the HCA 0A
   Port on the HCA 2

```
----------------- Specify HCA Attributes ------------


Specify or revise the values below.

Adapter ID of the HCA  . . 0A  +
Port on the HCA  . . . . . 2   +
```

Figure 8-16 HCD - Specify HCA Attributes: defining AID and port for 2827

9. Press Enter. HCD processes the normal Define Access and Candidate List panels for the new CHPID.
10. Complete the same procedure to define a new CIB type CHPID to the 2964 and press Enter. The example uses the following values:

    Channel path ID 90
    Channel path type CIB
    Operational mode SHR

11. Now enter the following values:

    Adapter ID of the HCA 08
    Port on the HCA 2

    Figure 8-17 shows the HCA attributes for the example.

```
----------------- Specify HCA Attributes -------


Specify or revise the values below.

Adapter ID of the HCA  . . 08  +
Port on the HCA  . . . . . 2   +
```

Figure 8-17 HCD - Specifying HCA attributes: defining AID and port for 2964

12. In HCD, display the Channel Path List for the 2964. Notice how HCD shows AID/P values (under CHID+) instead of PCHID values (Figure 8-18).

```
                              Channel Path List    Row 91 of 120 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+             Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
_ 86    27/1  CS5   SHR   __  __ __   Y   No  P Fram
_ 87    34/1  CS5   SHR   __  __ __   Y   No  P Fram
_ 90    08/2  CIB   SHR   __  __ __   N   No  ________________________________
_ 91    0A/2  CIB   SHR   __  __ __   Y   No  12x
_ 92    0C/2  CIB   SHR   __  __ __   N   No  12x
_ 93    0E/2  CIB   SHR   __  __ __   Y   No  12x
_ 94    0B/2  CIB   SHR   __  __ __   Y   No  1x
_ 95    0F/2  CIB   SHR   __  __ __   Y   No  1x
```

Figure 8-18 HCD - Channel Path List: displaying CIB CHPID and AID/P definition

13. Define a coupling link between the following CHPIDs:
   - 2827 CHPID 90 as type=CIB using AID/P=0A/2
   - 2964 CHPID 90 as type=CIB using AID/P=08/2=

14. Define a coupling link between those two CIB CHPIDs. Figure 8-19 shows the result.

```
------------------- View CF Control Units and Devices ---------


Peer CF side 1:
Processor ID . . . . . . . : SCZP401    Helix
Channel subsystem ID . . : 0
Channel path ID  . . . . : 90
Channel path type  . . . : CIB
Control unit number  . . :
Device number  . . . . . :
Number of devices  . . . :

Peer CF side 2:
Processor ID . . . . . . . : SCZP501    Sphinx
Channel subsystem ID . . : 0
Channel path ID  . . . . : 90
Channel path type  . . . : CIB
Control unit number  . . : FFF0
Device number  . . . . . : FF31
Number of devices  . . . : 7

ENTER to continue.
```

Figure 8-19 HCD - View CF Control Units and Devices: showing new CF link

15. From the HCD main menu, select option **1.3. Processor List** and press PF20 to display the SNA Address and Local Name (Figure 8-20).

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

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

                                     ----SNA Address + ---- ---Local---
/ Proc. ID Type +   Model +  Mode +  Network Name  CPC Name System Name
_ SCZP201  2097     E26      LPAR    USIBMSC       SCZP201  SCZP201
_ SCZP401  2827     H43      LPAR    USIBMSC       SCZP401  SCZP401
_ SCZP501  2964     N63      LPAR    USIBMSC       SCZP501  SCZP501
****************************** Bottom of data ********************************
```

Figure 8-20 HCD - Processor List: displaying the local system name

Tip: In most cases, the local system name value defaults to the CPC name, unless, during the addition of the new processor, the CPC name was not specified. In this case, HCD also leaves the local system name undefined.

System A (SCZP401) is known to itself as the LSYSTEM or local system, and it sees System B (SCZP501) as the CSYSTEM or connecting system.

System B (SCZP501) is known to itself as the LSYSTEM or local system, and it sees System A (SCZP401) as the CSYSTEM or connecting system.

CSYSTEM, and the other keywords that are associated with CIB type coupling links, can be seen in the IOCP deck. See Figure 8-50 on page 459.

Define any additional CHPIDs, control units and devices, FICON channel-to-channel connections (FCTCs), and so on, that you might be adding into the 2964 during the replace.

8.2.5 Defining more I/O

This section shows how to use HCD to define the following CHPIDs and PCIe functions:

- Defining a CF link over CHPID type CS5
- Defining a RoCE PCIe function
- Defining a zEDC EXPRESS PCIe function

Defining a CF link over CHPID type CS5

A new CHPID type called CS5 was introduced with the IBM z13 and uses a new PCIe-O SR ICA hardware card for the IBM z13 called *PCIe-O SR Integrated Coupling Adapter FC 0172*.

It runs at 8 GBps and up to a length of 150 meters. Presently its connectivity is supported only between one z13 and another z13.

The HCD definition process is similar to CHPIDs of type CIB except that the type CS5.

In this example, we define these items

- CHPID=0.86 to AID/Port ID=27/1 to processor=SCZP501
- CHPID=0.96 to AID/Port ID=27/2 to processor=SCZP502
- A CF link between these two CS5 CHPIDs

Complete these steps:

1. From the main HCD panel, select option **1.3. Processor List**. Enter **s** (work with attached channel paths) next to the first CPC (SCZP501) to which you want to define the CS5 CHPID and press Enter.
2. From the Channel Subsystem List, enter **s** (work with attached channel paths) next to the CSSID that you want. In our example, we use CSSID=0. Press Enter.
3. Enter **add** on the command line in the Channel Path List panel to add a new channel and then press Enter.
4. Enter the following values in the Add Channel Path panel (Figure 8-21) and press Enter:
   - Channel path ID: 86
   - Channel path type: CS5
   - Operational mode: SHR

```
-------------------------- Add Channel Path --------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 86    +            Channel ID  ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +
Description  . . . . . . CS5 CHPID definition___________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 8-21 HCD - Add Channel Path: defining a CS5 CHPID

5. Enter the following values in the Specify HCA Attributes panel (Figure 8-22) and then press Enter:
   - Adapter ID of the HCA: 27
   - Port on the HCA: 1

```
----------------- Specify HCA Attributes -------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 1   +
```

Figure 8-22 HCD - Specify HCA Attributes: defining CS5 HCA attributes

6. Select the required Access LPARs for CHPID access list. In our example, we use LPAR A0C (CF) and A01, A02, and A03 (OS). Press Enter. See Figure 8-23.

```
-------------------------- Define Access List ---------------------------
                                                              Row 1 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 86     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      OS
_ 0      A0B              B      CF/OS CHPID holder
/ 0      A0C              C      CF    CF7A
_ 0      A0D              D      CF    CF7B
_ 0      A0E              E      CF    CF7C
_ 0      A0F              F      CF    CF7D
/ 0      A01              1      OS    SC76
/ 0      A02              2      OS    SC74
/ 0      A03              3      OS    SC75
_ 0      A04              4      OS    ZZ0VM1
```

Figure 8-23 HCD - Define Access List: selecting partition for CS5 CHPID access

7. Select any Candidate LPARs for CHPID access list. In our example we will not select any candidate LPARs. Press Enter.
8. From the main HCD panel, select option **1.3. Processor List**. Enter `s` (work with attached channel paths) next to the second CPC (SCZP502) to which you want to define the CS5 CHPID, and then press Enter.
9. From the Channel Subsystem List, enter `s` (work with attached channel paths) next to the CSSID. In our example, we use CSSID=0. Press Enter.
10. Enter `add` on the command line in the Channel Path List panel to add a new channel and then press Enter.

11. Enter the following values in the Add Channel Path panel (Figure 8-24) and press Enter.
   – Channel path ID: 96
   – Channel path type: CS5
   – Operational mode: SHR

```
--------------------------- Add Channel Path ------------------------


Specify or revise the following values.

Processor ID . . . . : SCZP502      test machine definition
Configuration mode . : LPAR
Channel Subsystem ID : 0

Channel path ID  . . . . 96    +           Channel ID  ___  +
Number of CHPIDs . . . . 1
Channel path type  . . . CS5   +
Operation mode . . . . . SHR   +
Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +
Description  . . . . . . CS5 CHPID definition____________

Specify the following values only if connected to a switch:
Dynamic entry switch ID  __  + (00 - FF)
Entry switch ID  . . . . __  +
Entry port . . . . . . . __  +
```

Figure 8-24   HCD - Add Channel Path: defining a CS5 CHPID

12. Enter the following values in the Specify HCA Attributes panel (Figure 8-25) and then press Enter:
   – Adapter ID of the HCA: 27
   – Port on the HCA: 2

```
----------------- Specify HCA Attributes ------


Specify or revise the values below.

Adapter ID of the HCA  . . 27  +
Port on the HCA  . . . . . 2   +
```

Figure 8-25   HCD - Specify HCA Attributes: defining CS5 HCA attributes

13. Select the required Access LPAR for CHPID access list. In our example, we use LPAR A0D (CF) and A01, A02, and A03 (OS). Then press Enter. See Figure 8-26.

```
-------------------------- Define Access List ---------------------------
                                                               Row 1 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 96     Channel path type  . : CS5
Operation mode . . . : SHR    Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0A              A      CF/OS
_ 0      A0B              B      CF/OS
_ 0      A0C              C      CF
/ 0      A0D              D      CF
_ 0      A0E              E      CF/OS
_ 0      A0F              F      CF/OS
/ 0      A01              1      OS
/ 0      A02              2      OS
/ 0      A03              3      OS
_ 0      A04              4      CF/OS
```

Figure 8-26 HCD - Define Access List: selecting partition for CS5 CHPID access

14. Select any Candidate LPARs for CHPID access list. In our example, we do not select any candidate LPARs. Press Enter.

15. Return to the first CHPID definition (0.86 on processor SCZP501) and type **f** (connect CF channel paths) next to CHPID 0.86 and press Enter. See Figure 8-27.

```
                              Channel Path List    Row 91 of 117 More:      >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP501      Sphinx
Configuration mode . : LPAR
Channel Subsystem ID : 0

        CHID+            Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
f 86    27/1  CS5   SHR   __  __ __   N   No  CS5 CHPID definition
_ 87    34/1  CS5   SHR   __  __ __   N   No  P Fram
_ 90    08/2  CIB   SHR   __  __ __   Y   No  12x P201 T/O
_ 94    0B/2  CIB   SHR   __  __ __   Y   No  1x
_ 95    0F/2  CIB   SHR   __  __ __   Y   No  1x
```

Figure 8-27 HCD - Channel Path List: connecting a CF channel path

16.Type **p** (CF channel path connectivity list) next to CHPID 0.86 and press Enter. See Figure 8-28.

```
                     CF Channel Path Connectivity List             Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N     SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N     SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N     SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N     SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N     SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N     SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
p 86  27/1  Y  CS5  SHR  N
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N     SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 8-28 HCD - CF Channel Path Connectivity List: creating a CF link

17.Enter the following values in the Connect to CF Channel Path panel (Figure 8-29) and then press Enter:

– Destination processor ID: SCZP502
– Destination channel subsystem ID: 0
– Destination channel path ID: 96

```
--------------- Connect to CF Channel Path --------


Specify the following values.

Source processor ID  . . . . . : SCZP501
Source channel subsystem ID  . : 0
Source channel path ID . . . . : 86
Source channel path type . . . : CS5

Destination processor ID . . . . . . SCZP502_  +
Destination channel subsystem ID . . 0   +
Destination channel path ID  . . . . 96  +

Timing-only link . . . . . . . . . . No
```

Figure 8-29 HCD - Connect to CF Channel Path: specifying the destination attributes

18. Accept or override the control unit and device numbers for processor *SCZP501*, then press Enter. See Figure 8-30.

```
---------------- 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 . . . . . . . : SCZP501
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 86        Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF6  +

Device number  . . . . . . . FE3D
Number of devices  . . . . . 8
```

Figure 8-30 HCD - Add CF Control Unit and Devices: accept or override default values

19. Accept or override the Control Unit and Device numbers for processor *SCZP502*, then press Enter. See Figure 8-31.

```
---------------- 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 . . . . . . . : SCZP502
Channel subsystem ID . . . : 0
Channel path ID  . . . . . : 96        Operation mode . . : SHR
Channel path type  . . . . : CS5

Control unit number  . . . . FFF5  +

Device number  . . . . . . . FE2D
Number of devices  . . . . . 8
```

Figure 8-31 HCD - Add CF Control Unit and Devices: accept or override default values

HCD returns to the CF Channel Path Connectivity List panel (Figure 8-32) where you can see the CF link now defined.

```
                    CF Channel Path Connectivity List            Row 1 of 17
Command ===> ____________________________________________ Scroll ===> CSR

Select one or more channel paths, then press Enter.

Source processor ID . . . . . : SCZP501   Sphinx
Source channel subsystem ID . : 0
Source partition name . . . . : *

  ----------Source----------   -----------Destination-----------   -CU-  -#-
/ CHP CHID  CF Type Mode Occ   Proc.CSSID CHP CHID  CF Type Mode   Type  Dev
_ 80  08/1  Y  CIB  SHR  N     SCZP501.0  82  0C/1  Y  CIB  SHR    CFP   7
_ 81  0A/1  Y  CIB  SHR  N     SCZP501.0  83  0E/1  Y  CIB  SHR    CFP   7
_ 82  0C/1  Y  CIB  SHR  N     SCZP501.0  80  08/1  Y  CIB  SHR    CFP   7
_ 83  0E/1  Y  CIB  SHR  N     SCZP501.0  81  0A/1  Y  CIB  SHR    CFP   7
_ 84  0B/1  Y  CIB  SHR  N     SCZP501.0  94  0B/2  Y  CIB  SHR    CFP   32
_ 85  0F/1  Y  CIB  SHR  N     SCZP501.0  95  0F/2  Y  CIB  SHR    CFP   32
_ 86  27/1  Y  CS5  SHR  N     SCZP502.0  96  27/2  Y  CS5  SHR    CFP   8
_ 87  34/1  Y  CS5  SHR  N
_ 90  08/2  Y  CIB  SHR  N     SCZP201.1  BA  0B/2  Y  CIB  SHR    STP
```

Figure 8-32 HCD - CF Channel Path Connectivity List: CF link now created

Defining a RoCE PCIe function

RDMA over Converged Ethernet (RoCE) uses a new PCIe hardware card for the zEC12 and IBM z13 and is named *10GbE RoCE Express SR* (FC 0411).

To use RoCE hardware functions, Function IDs, Virtual Function IDs, and Physical Network IDs must be defined in HCD to the PCHID that has been assigned to the RoCE hardware cards 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 RoCE PCHIDs.

These steps describe only the definition process. However, if you want deeper understanding of how RoCe works, see Chapter 14, “Configure Open System Adaptor (OSA) and RoCE Express” on page 653.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 8-2 on page 452)

Example 8-1 PCHID Report - RoCE information from PCHID report

```
Source          Cage  Slot  F/C   PCHID/Ports or AID           Comment
19/02/J01       Z22B  20    0411  140                     RG1  NEW
19/12/J01       Z08B  03    0411  208                     RG2  NEW
```

In this example, we define these items:

- PCHID=140 to Function IDs 000 and 001 on SCZP502
- PCHID=208 to Function IDs 010 and 011 on SCZP502

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 (SCZP502) to which you want to define the RoCE functions and press Enter. See Figure 8-40 on page 453.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
******************************* Bottom of data ********************************
```

Figure 8-33 HCD - Processor List: adding PCIe functions to a processor

2. Type **add** on the command line in the PCIe Function List panel (Figure 8-41 on page 453) to add a new PCIe function.

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 8-34 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 8-35) and press Enter.
   - Function ID: `000`
   - Type: `ROCE`
   - PCHID: `140`
   - Virtual Function ID: `1`
   - Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID . . . . : SCZP502      test machine definition

Function ID . . . . . . 000
Type  . . . . . . . . . ROCE          +

PCHID . . . . . . . . . 140
Virtual Function ID . . 1_  +

Description . . . . . . RoCE F=000 VF=1 PCHID=140_______
```

Figure 8-35 HCD - PCIe Function List: adding PCIe functions to a processor

4. Enter the following value in the Add/Modify Physical Network IDs panel (Figure 8-36) and press Enter.
   - Physical network ID 1 = `ITSOPNET1`

```
-------------------- Add/Modify Physical Network IDs ------------------


If the PCHID is associated to one or more physical networks, specify
each physical network ID corresponding to each applicable physical port

Physical network ID 1  . . ITSOPNET1_______
Physical network ID 2  . . ________________
Physical network ID 3  . . ________________
Physical network ID 4  . . ________________
```

Figure 8-36 HCD - Add/Modify Physical Network IDs: adding network ID

5. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Then press Enter. See Figure 8-37.

```
-------------------------- Define Access List ----------------------------
                                                              Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 8-37   HCD - Define Access List: selecting partition for Function access

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 (Figure 8-38) where you can see the Function now defined.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
******************************* Bottom of data ********************************
```

Figure 8-38   HCD - PCIe Function List: Function now created

7. Define the other Function IDs as per the example so far. See Figure 8-39.

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

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE F=000 VF=1 PCHID=140
_ 001    140    2    ROCE            RoCE F=000 VF=2 PCHID=140
_ 010    208    1    ROCE            RoCE F=000 VF=1 PCHID=208
_ 011    208    2    ROCE            RoCE F=000 VF=2 PCHID=208
******************************* Bottom of data ********************************
```

Figure 8-39 HCD - PCIe Function List: All functions now created

Defining a zEDC EXPRESS PCIe function

The zEnterprise Data Compression (zEDC) uses a new PCIe hardware card for the zEC12 and IBM z13 named *zEDC Express FC 0420*.

To use zEDC hardware functions, Function IDs and Virtual Function IDs must be defined in HCD to the PCHID that has been assigned to the zEDC hardware cards 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.

Here we show only the definition process. If you want a deeper understanding of how zEDC works, see *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

First, for information about the PCHID and Resource Group (RG), see the PCHID report for the processor (Example 8-2).

Example 8-2 PCHID Report - zEDC information from PCHID report

```
Source          Cage  Slot  F/C    PCHID/Ports or AID                Comment
15/14/J01       Z15B  19    0420   1BC                        RG1    NEW
15/05/J01       Z08B  38    0420   27C                        RG2    NEW
```

In this example, we define these items:

- PCHID=1BC to Function IDs 020, 021, and 022 on SCZP502
- PCHID=27C to Function IDs 030, 031, and 032 on SCZP502

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 (SCZP502) to which you want to define the zEDC functions and press Enter. See Figure 8-40.

```
                                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
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
_ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
f SCZP502  2964     N63      LPAR  08DA882964 test machine definition
******************************* Bottom of data ********************************
```

Figure 8-40 HCD - Processor List: adding PCIe functions to a processor

2. Type **add** on the command line in the PCIe Function List panel (Figure 8-41) to add a new PCIe function.

```
                              PCIe Function List
Command ===> add____________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 8-41 HCD - PCIe Function List: adding PCIe functions to a processor

3. Enter the following values in the Add PCIe Function panel (Figure 8-42) and press Enter.
   – Function ID: `020`
   – Type: `ZEDC-EXPRESS`
   – PCHID: `1BC`
   – Virtual Function ID: `1`
   – Description: `optional`

```
--------------------------- Add PCIe Function -----------------------


Specify or revise the following values.

Processor ID  . . . . : SCZP502      test machine definition

Function ID . . . . . . 020
Type  . . . . . . . . . ZEDC-EXPRESS  +

PCHID . . . . . . . . . 1BC
Virtual Function ID . . 1   +

Description . . . . . . zEDC F=020 VF=1 PCHID=1BC_______
```

Figure 8-42 HCD - PCIe Function List: adding PCIe functions to a processor

4. Select the required Access LPAR for Function access list. In our example, we use LPAR A01 (OS). Then, press Enter. See Figure 8-43.

```
-------------------------- Define Access List ----------------------------
                                                           Row 7 of 15
Command ===> _________________________________________ Scroll ===> CSR

Select one or more partitions for inclusion in the access list.

Function ID  . . . . : 020

/ CSS ID Partition Name   Number Usage Description
/ 0      A01              1      OS
_ 0      A02              2      OS
_ 0      A03              3      OS
_ 0      A04              4      CF/OS
_ 0      A05              5      CF/OS
_ 0      A06              6      CF/OS
_ 0      A07              7      OS
_ 0      A08              8      OS
_ 0      A09              9      CF/OS
**************************** Bottom of data *****************************
```

Figure 8-43 HCD - Define Access List: selecting partition for Function access

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 (Figure 8-44) where you can see that the Function is now defined.

```
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
******************************* Bottom of data ********************************
```

Figure 8-44 HCD - PCIe Function List: Function now created

6. Now define the other Function IDs according to this example. See Figure 8-45.

```
                              PCIe Function List      Row 1 of 6 More:       >
Command ===> _______________________________________________ Scroll ===> CSR

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

Processor ID . . . . : SCZP502      test machine definition



/ FID    PCHID  VF+  Type+           Description
_ 020    1BC    1    ZEDC-EXPRESS    zEDC F=020 VF=1 PCHID=1BC
_ 021    1BC    2    ZEDC-EXPRESS    zEDC F=021 VF=2 PCHID=1BC
_ 022    1BC    3    ZEDC-EXPRESS    zEDC F=022 VF=3 PCHID=1BC
_ 030    27C    1    ZEDC-EXPRESS    zEDC F=030 VF=1 PCHID=27C
_ 031    27C    2    ZEDC-EXPRESS    zEDC F=031 VF=2 PCHID=27C
_ 032    27C    3    ZEDC-EXPRESS    zEDC F=032 VF=3 PCHID=27C
******************************* Bottom of data ********************************
```

Figure 8-45 HCD - PCIe Function List: All Functions now created

8.3 HCD: Validating the 2964 work IODF

To validate the 2964 work IODF by using HCD, complete the following steps:

1. Select HCD option **2.12. Build validated work I/O definition file**.
2. Review the message list and correct any errors.
3. Press PF3 to continue. You receive the following message:

   `Requested action successfully processed`

4. Select HCD option **6.4. View I/O Definition File Information**. The IODF type is now indicated as `Work - Validated` (Figure 8-46).

```
-------------- View I/O Definition File Information ------



IODF name  . . . . . . : 'SYS6.IODFB8.WORK'
IODF type  . . . . . . : Work - Validated
IODF version . . . . . : 5

Creation date  . . . . : 2014-11-10
Last update  . . . . . : 2014-11-10  15:34

Volume serial number . : IODFBK
Allocated space  . . . : 4000    (Number of 4K blocks)
Used space . . . . . . : 1983    (Number of 4K blocks)
   thereof utilized (%)  98
Activity logging . . . : No
Multi-user access  . . : No
Backup IODF name . . . :

Description  . . . . . :
```

Figure 8-46 HCD - View I/O Definition File Information: validated work IODF

8.4 Creating the IOCP statements for the CHPID Mapping Tool

To create the IOCP file for the CHPID Mapping Tool, complete the following steps:

Tip: You might prefer to use HCM to create the IOCP statements file and transfer the file to your workstation. You can then start the CHPID Mapping Tool, create an updated IOCP statements file, and transfer the file back to the host.

1. Select HCD option **3. Build IOCP input data set** and press Enter (Figure 8-47).

```
----- 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
```

Figure 8-47 HCD - Activate or Process Configuration Data: build IOCP for SCZP501

2. HCD displays the list of available processors (Figure 8-48). Select the SCZP501 processor by typing a forward slash (/) next to it and pressing Enter.

```
------------------------- Available Processors ---------------------------
                                                               Row 1 of 3
Command ===> ____________________________________________________________

Select one.

  Processor ID  Type     Model    Mode  Description
  SCZP201       2097     E26      LPAR  Eclipse
  SCZP401       2827     H43      LPAR  Helix
/ SCZP501       2964     N63      LPAR  Sphinx
**************************** Bottom of data *****************************
```

Figure 8-48 HCD - 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 (Figure 8-49).

   Complete the following fields:

   – Title1: `IODFB8`
   – IOCP input data set: `'SYS6.IODFB8.IOCPIN.SCZP501'`
   – Input to Stand-alone IOCP: `Yes`
   – The Job statement information: Complete this information for your installation.

```
------------------------ Build IOCP Input Data Set -----------------------


Specify or revise the following values.

IODF name . . . . . . . . . : 'SYS6.IODFB8.WORK'
Processor ID  . . . . . . . : SCZP501
Title1 . IODFB8__________________________________________________________
Title2 : SYS6.IODFB8.WORK - 2014-11-10 15:34

IOCP input data set
'SYS6.IODFB8.IOCPIN.SCZP501'____________________________
Input to Stand-alone IOCP?  Yes  (Yes or No)

Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 8-49   HCD - 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 (Figure 8-50). This data set is used as input into the CHPID Mapping Tool. Also shown here is the LSYSTEM keyword and the CHPID statement that shows the CPATH and CSYSTEM keywords.

```
ID     MSG1='IODFB8',                                          *
       MSG2='SYS6.IODFB8.WORK - 2014-11-10 15:34',             *
       SYSTEM=(2964,1),LSYSTEM=SCZP501,                        *
       TOK=('SCZP501',00800001DA8729641534090801l4314F00000000,*
       00000000,'14-11-10','15:34:09','........','........')
RESOURCE PARTITION=((CSS(0),(A0A,A),(A0B,B),(A0C,C),(A0D,D),(A*
       0E,E),(A0F,F),(A01,1),(A02,2),(A03,3),(A04,4),(A05,5),(A*
       06,6),(A07,7),(A08,8),(A09,9)),(CSS(1),(*,1),(*,2),(*,3)*
       ,(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*
       *,D),(*,E),(*,F)),(CSS(2),(*,1),(*,2),(*,3),(*,4),(*,5),*
       (*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(**
       ,F)),(CSS(3),(*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*
       *,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(4),*
       (*,1),(*,2),(*,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(**
       ,A),(*,B),(*,C),(*,D),(*,E),(*,F)),(CSS(5),(*,1),(*,2),(*
       *,3),(*,4),(*,5),(*,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,*
       C),(*,D),(*,E),(*,F)))
FUNCTION FID=000,VF=1,PCHID=140,PNETID=ITSOPNET1,              *
       PART=((A01),(=)),TYPE=ROCE
FUNCTION FID=010,VF=1,PCHID=208,PNETID=ITSOPNET1,              *
       PART=((A02),(=)),TYPE=ROCE
FUNCTION FID=011,VF=2,PCHID=208,PNETID=ITSOPNET1,              *
       PART=((A03),(=)),TYPE=ROCE
...........................................................
CHPID PATH=(CSS(0),90),SHARED,PARTITION=((A0A,A0B,A09),(=)),   *
       CPATH=(CSS(0),90),CSYSTEM=SCZP401,AID=0A,PORT=2,TYPE=CIB
...........................................................
CNTLUNIT CUNUMBR=FFFC,PATH=((CSS(0),80,90)),UNIT=CFP
IODEVICE ADDRESS=(FFD6,007),CUNUMBR=(FFFC),UNIT=CFP
IODEVICE ADDRESS=(FFF2,007),CUNUMBR=(FFFC),UNIT=CFP
```

Figure 8-50 HCD - IOCP input data set: contents (truncated)

Part of the `TOK` statement is now replaced with dots (Example 8-3).

Example 8-3 HCD - IOCP file: TOK statement

```
TOK=('SCZP501',00800001DA8729641534090801l4314F00000000,*
00000000,'14-11-10','15:34:09','........','........')
```

These dots ensure that this IOCP file cannot be written to a processor and used for a power-on reset. This IOCP file was created from a validated work IODF and not a production IODF, which is something that can be done only for processors that contain PCHID definitions.

> **Important:** When an IOCP statement file is exported from a validated work IODF 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 file from TSO to your workstation. Use a workstation file transfer facility such as the one in 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 `SCZP501in.iocp`.

8.5 CMT: Assigning PCHIDs to CHPIDs

The following steps use the output from the previous set of HCD steps (IOCP) and the 2964 order process (CFReport). Use the CHPID Mapping Tool (CMT) to assign PCHIDs to each of the CHPIDs for the 2964.

For this process, the CMT must be downloaded. For more information about downloading and installing the CMT, see 3.1.5, "CHPID Mapping Tool (CMT)" on page 44. If CMT is already installed, verify that the latest updates are installed. For more information, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

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 PCHID conflict.
4. Process the hardware resolution.
5. Resolve manually the CIB 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 and transfer it back to the host z/OS image.

> **Requirement:** When you upgrade from a 2817 or a 2827 to a 2964, you must use the CHPID Mapping Tool level that supports the 2964.

8.5.1 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. Right-click the Projects pane and select **New** → **Standard CMT Project** (Figure 8-51).

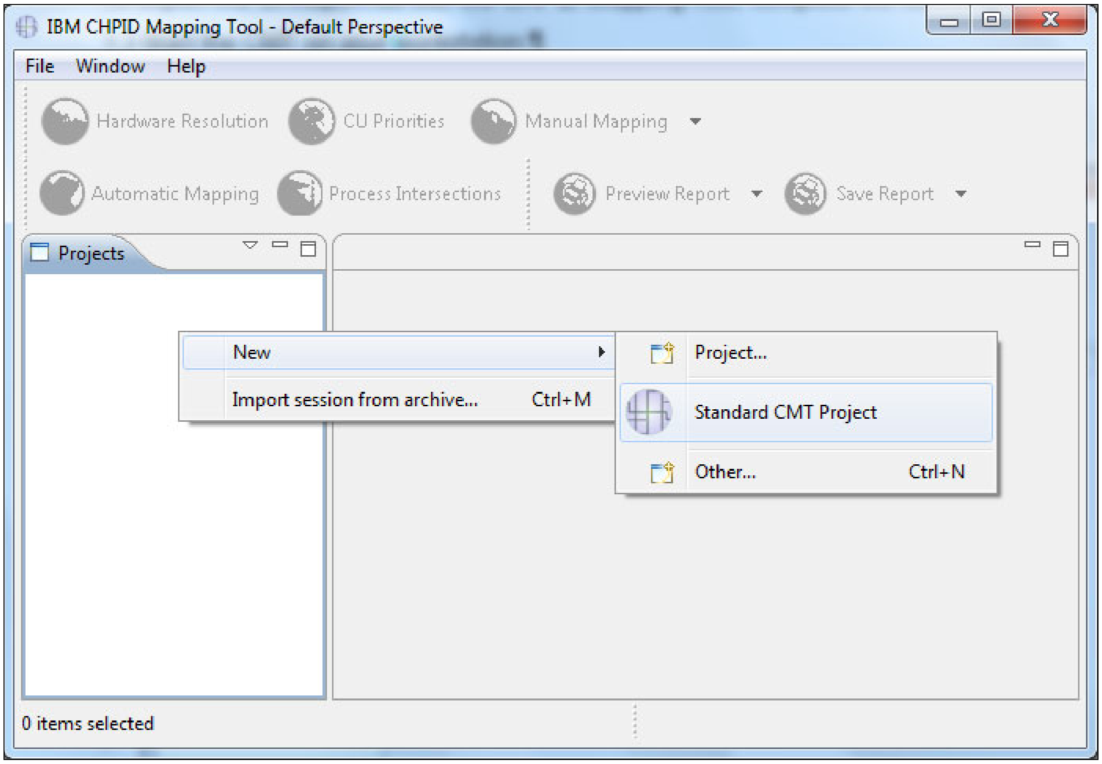

Figure 8-51 CMT - Creating a new CHPID Mapping Tool Project

3. The New CHPID Mapping Tool Project window opens (Figure 8-52). Use this window to specify a project name and then click **Next**.

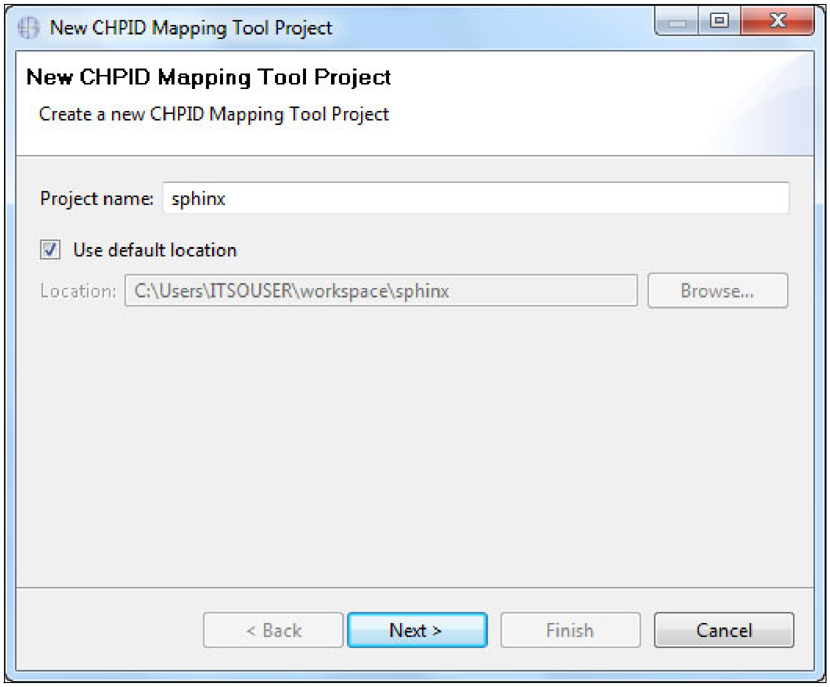

Figure 8-52 CMT - New CHPID Mapping Tool Project name

4. Import the `CFReport` file into the CHPID Mapping Tool by specifying the name in the CFReport file field, and then click **Finish** (Figure 8-53).

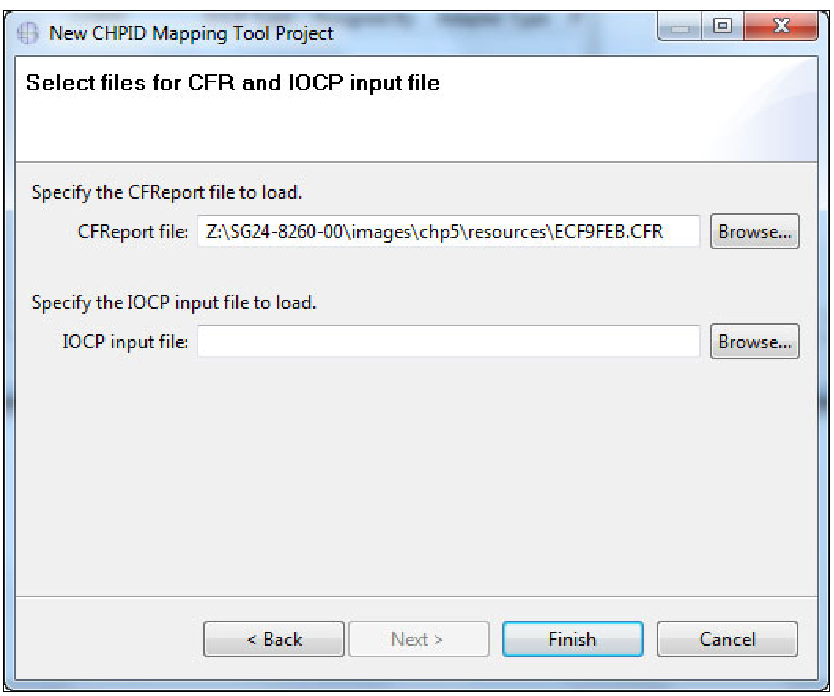

Figure 8-53 CMT - Specifying the CFReport file

If you click **Finish** but did not select an IOCP file, you receive the message shown in Figure 8-54.

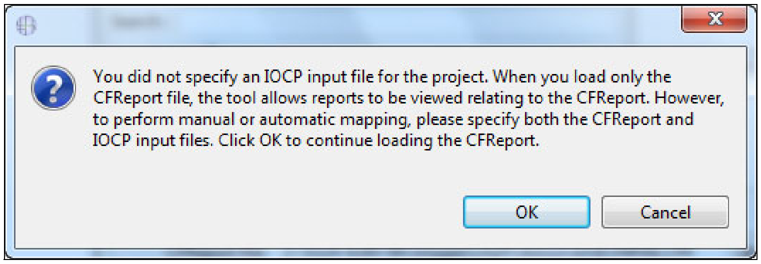

Figure 8-54   CMT - Warning message for not specifying an IOCP file

A window shows the progress of reading the CFReport file (Figure 8-55).

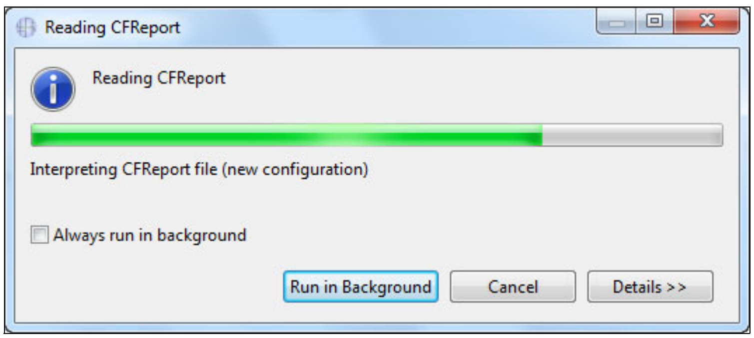

Figure 8-55   CMT - Reading the CFReport file

The information from the CFReport file is displayed on the Hardware pane (Figure 8-56).

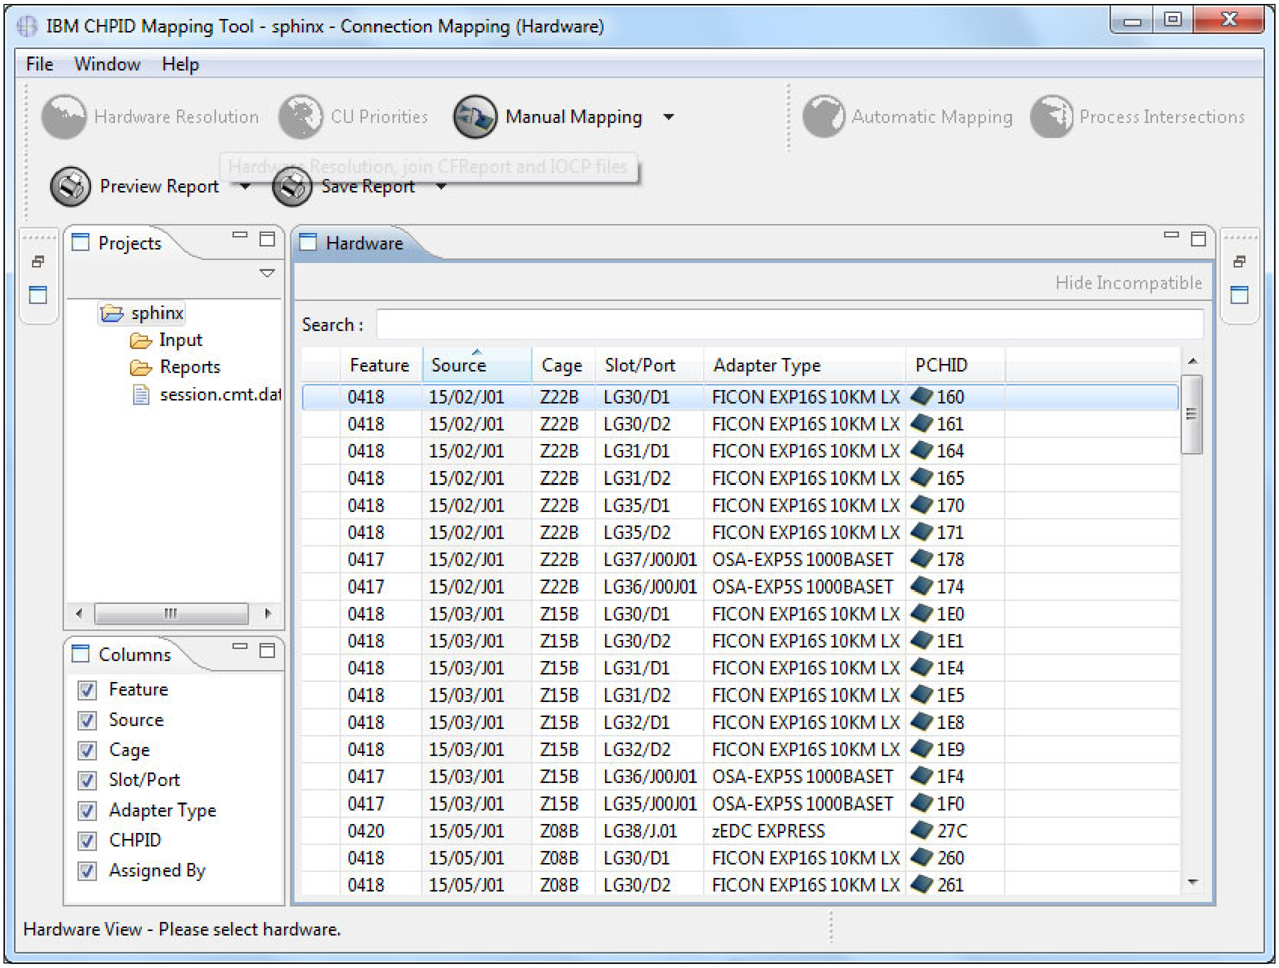

Figure 8-56 CMT - Imported CFReport file

8.5.2 Importing the 2964 IOCP file into the CHPID Mapping Tool

To import the 2964 IOCP file into the CHPID Mapping Tool, complete the following steps:

1. Right-click in the Projects pane and select **Import IOCP input file** (Figure 8-57).

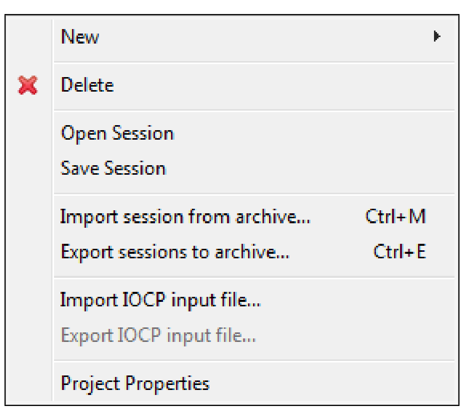

Figure 8-57 CMT - Importing the IOCP file

2. Select the IOCP file on your workstation to import into the CHPID Mapping Tool and click **Finish** (Figure 8-58).

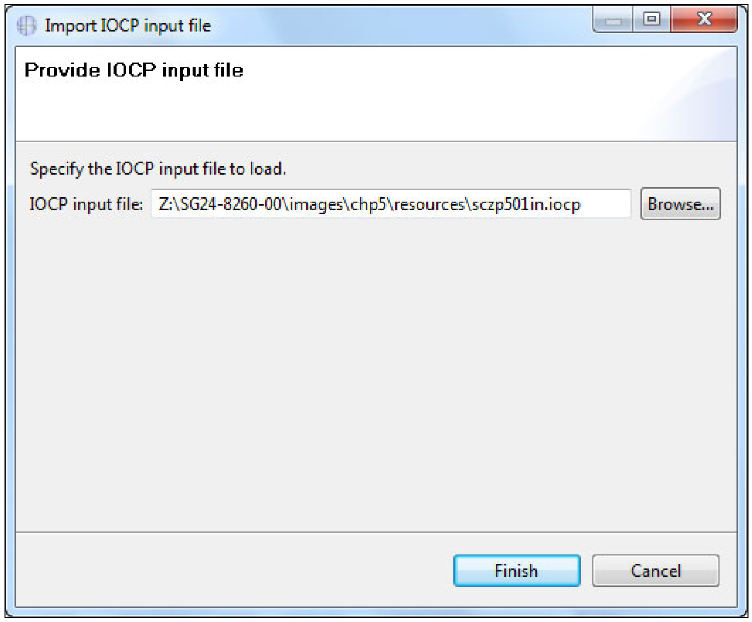

Figure 8-58 CMT - Specifying the IOCP file for import

3. In the Projects pane, under the Input tab, expand the IOCP tab, right-click the IOCP file, and select **Read Selected IOCP (**Figure 8-59).

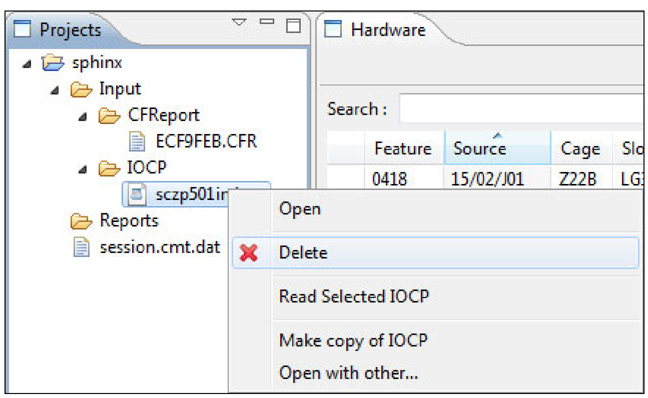

Figure 8-59 CMT - Reading the selected IOCP

A window displays the progress information (Figure 8-60).

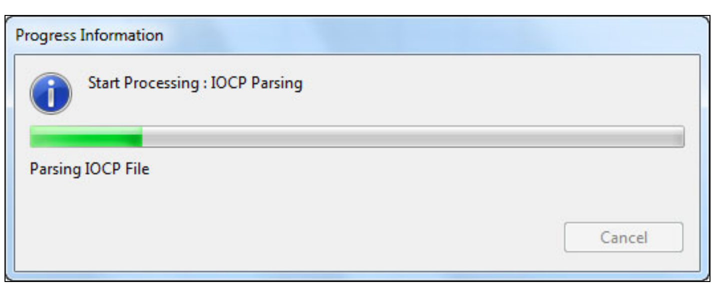

Figure 8-60 CMT - Processing the IOCP file

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 (Figure 8-61) includes three tabbed panes: Projects, Hardware Resolution and Adapter Type Summary. Hardware Resolution is the middle pane and the Adapter Type Summary is on the right.

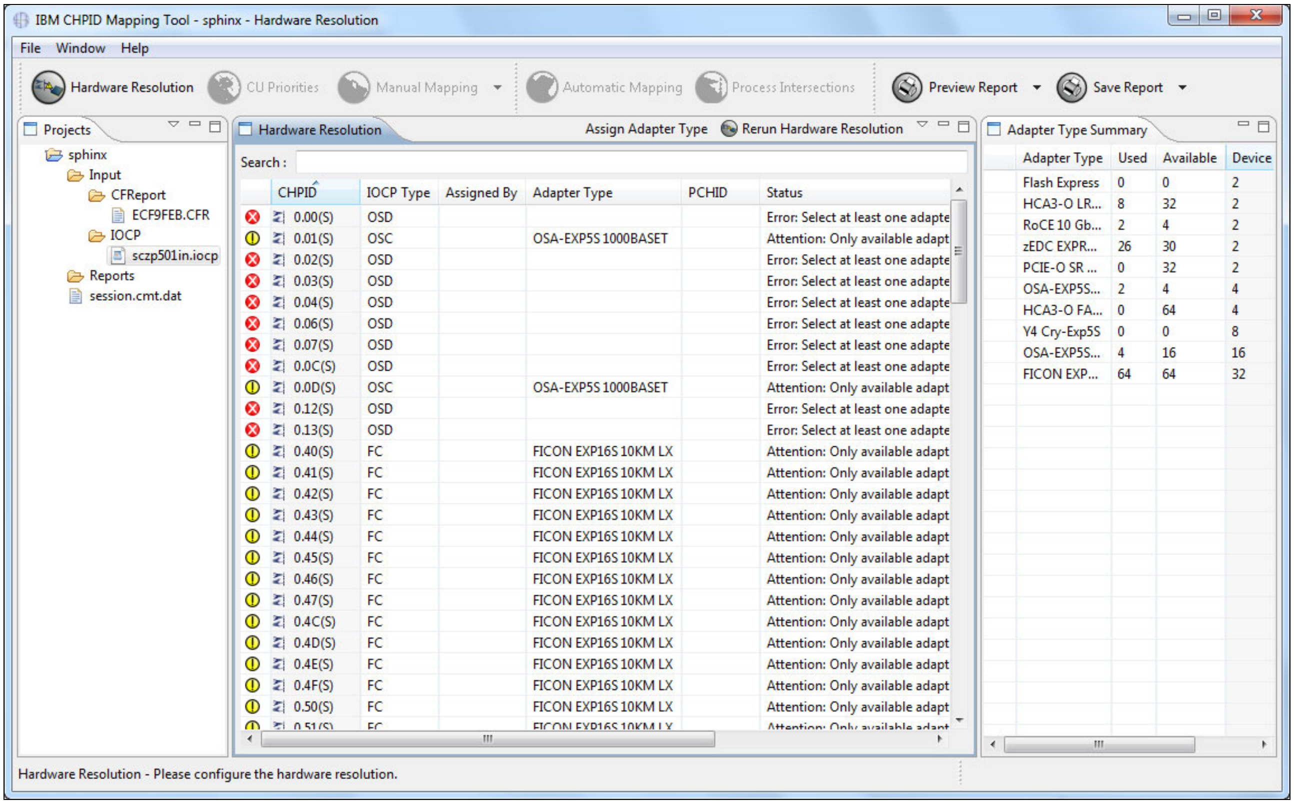

Figure 8-61 CMT - Hardware Resolution after Imported CFReport file

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 issues the following output:

- Hardware Resolution: This window lists all CHPIDs that were found; the Status column shows the CHPID information to be investigated. In the example, investigate the status; Table 8-3 lists status messages and possible resolutions.

Table 8-3 Status messages and possible resolutions

Status	Explanation	Resolution (if required)
No hardware found	AID values or PCHID values are present that are not found in the hardware. This situation could occur when you are replacing hardware for an MES and the IOCP file contains a PCHID value for the old hardware. (The IOCP file contains a PCHID value for the hardware being removed.)	If you have any CHPIDs of IOCP type CIB, 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. Do the following steps for CIB CHPIDs: 1. Remove the AID values. 2. Do one of the following tasks: - Inside the CHPID Mapping Tool, perform manual mapping to associate these CHPIDs with AIDs. - Assign the AID values outside the tool, for example, using Hardware Configuration Definition (HCD). 3. Replace the IOCP file.
Select at least one adapter type.	A adapter type is not assigned to the current row.	Assign a adapter type to IOCP type.
Adapter_type is not compatible with *IOCP_type*.	Adapter type assigned for the CHPID is not compatible with the IOCP type specified by the IOCP file.	See Performing hardware resolution for a type mismatch
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 need to change the IOCP file or obtain additional hardware
PCHID_1 moved to new channel ID: PCHID_2 **Example**: 520 moved to 1E2	You are replacing hardware for an MES, and the IOCP file contains a PCHID value for the old hardware, which is being removed. This PCHID value has moved from an old machine to the PCHID value for the new hardware. PCHID_1 is the first PCHID value (for example, 520) and PCHID_2 is the second PCHID 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 if you prefer a different assignment.

- Manual mapping CIB CHPIDs: Availability Mapping cannot be used until all CIB CHPIDs are resolved. You can use manual mapping to resolve any CIB 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 PCHID in the Manual window. If this option is not selected (it has no check mark), then availability PCHIDs 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 might require recabling after availability assignments.

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

    Regardless of whether the channels are moving or not, the CHPID Mapping Tool either assigns PCHIDs 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 PCHID where the affected CHPIDs are assigned) before you use the CHPIDs.

    By default, this option is not selected.

  If none of these options is selected, availability works only on CHPIDs that do not have PCHIDs assigned.

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

  > **Attention:** If you run **Reset CHPIDs assigned by IOCP**, it will reset any previously mapped CHPID assignments and 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 prior CHPID-to-PCHID relationship does not cause unacceptable availability.

8.5.3 Resolving CHPIDs with PCHID conflict

The CMT displays the CHPIDs with PCHID conflicts (Figure 8-62).

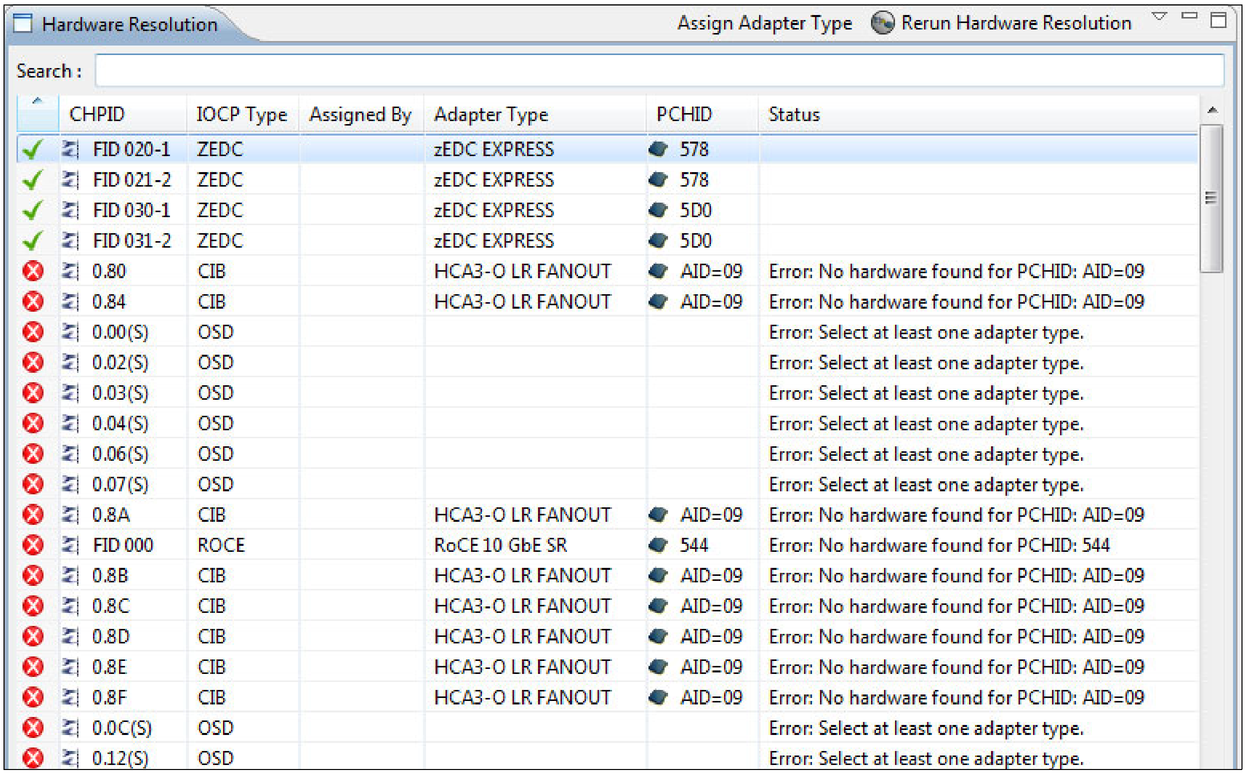

Hardware Resolution   Assign Adapter Type   Rerun Hardware Resolution

Search :

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 020-1	ZEDC		zEDC EXPRESS	578	
FID 021-2	ZEDC		zEDC EXPRESS	578	
FID 030-1	ZEDC		zEDC EXPRESS	5D0	
FID 031-2	ZEDC		zEDC EXPRESS	5D0	
0.80	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.84	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.00(S)	OSD				Error: Select at least one adapter type.
0.02(S)	OSD				Error: Select at least one adapter type.
0.03(S)	OSD				Error: Select at least one adapter type.
0.04(S)	OSD				Error: Select at least one adapter type.
0.06(S)	OSD				Error: Select at least one adapter type.
0.07(S)	OSD				Error: Select at least one adapter type.
0.8A	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
FID 000	ROCE		RoCE 10 GbE SR	544	Error: No hardware found for PCHID: 544
0.8B	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8C	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8D	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8E	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.8F	CIB		HCA3-O LR FANOUT	AID=09	Error: No hardware found for PCHID: AID=09
0.0C(S)	OSD				Error: Select at least one adapter type.
0.12(S)	OSD				Error: Select at least one adapter type.

Figure 8-62 CMT - CHPIDs with PCHID conflicts

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

- An X in a red circle: This indicates an error.
- A green check mark: This indicates that the tool successfully resolved the specified Channel Type.

The example has these reasons to resolve hardware resolution issues:

- The PCHID channel type changed.
- Defined PCHID is not compatible with the channel path at a particular location.
- There are enough ports in the hardware.
- A type mismatch exists between a CHPID and its associated channel type.

8.5.4 Hardware Resolution

In the example, the CHPID Mapping Tool displays an X in the first column of the Hardware Resolution pane (Figure 8-63) and is related to these error types: `No hardware found` and `FICON EXP8S 10KM LX is not compatible with OSD`.

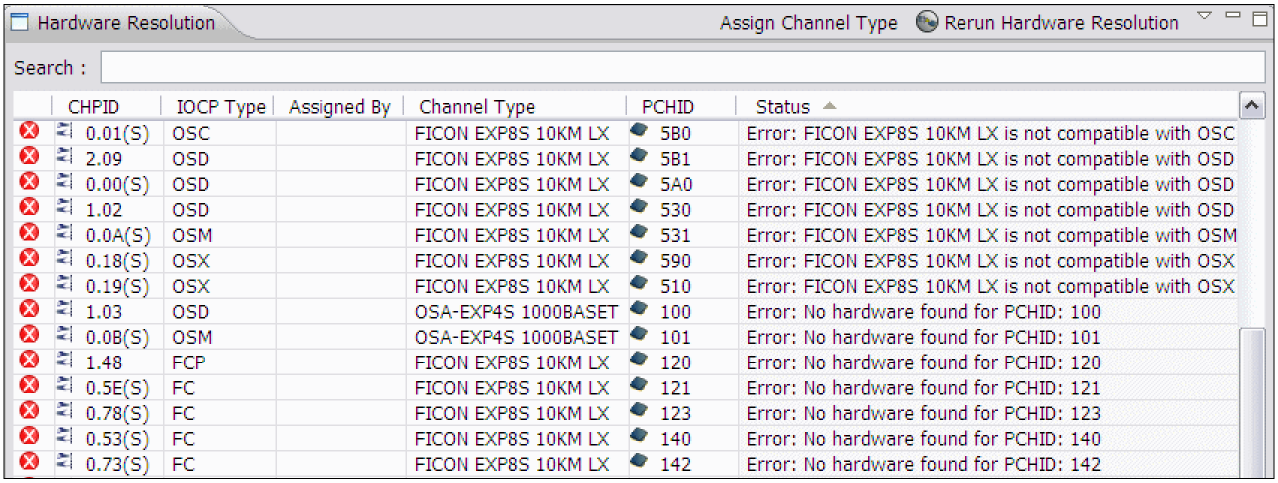

CHPID	IOCP Type	Assigned By	Channel Type	PCHID	Status
0.01(S)	OSC		FICON EXP8S 10KM LX	5B0	Error: FICON EXP8S 10KM LX is not compatible with OSC
2.09	OSD		FICON EXP8S 10KM LX	5B1	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.00(S)	OSD		FICON EXP8S 10KM LX	5A0	Error: FICON EXP8S 10KM LX is not compatible with OSD
1.02	OSD		FICON EXP8S 10KM LX	530	Error: FICON EXP8S 10KM LX is not compatible with OSD
0.0A(S)	OSM		FICON EXP8S 10KM LX	531	Error: FICON EXP8S 10KM LX is not compatible with OSM
0.18(S)	OSX		FICON EXP8S 10KM LX	590	Error: FICON EXP8S 10KM LX is not compatible with OSX
0.19(S)	OSX		FICON EXP8S 10KM LX	510	Error: FICON EXP8S 10KM LX is not compatible with OSX
1.03	OSD		OSA-EXP4S 1000BASET	100	Error: No hardware found for PCHID: 100
0.0B(S)	OSM		OSA-EXP4S 1000BASET	101	Error: No hardware found for PCHID: 101
1.48	FCP		FICON EXP8S 10KM LX	120	Error: No hardware found for PCHID: 120
0.5E(S)	FC		FICON EXP8S 10KM LX	121	Error: No hardware found for PCHID: 121
0.78(S)	FC		FICON EXP8S 10KM LX	123	Error: No hardware found for PCHID: 123
0.53(S)	FC		FICON EXP8S 10KM LX	140	Error: No hardware found for PCHID: 140
0.73(S)	FC		FICON EXP8S 10KM LX	142	Error: No hardware found for PCHID: 142

Figure 8-63 CMT - Hardware resolution status errors

> **More information:** For more information about these error messages, see the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

The options that must be reset are as follows:

- Reset Incompatible (Hardware - I/O) Entries: (shown in example)
- Reset "Error: No hardware found" Entries: (shown in example)
- Reset "Select at least one adapter type": (shown in example)
- Reset "Required hardware for type IOCP_type not available": (*not* shown in example)
- Reset "PCHID_1 moved to new channel ID: PCHID_2": (*not* shown in example)

Reset Incompatible (Hardware - I/O) Entries

The Channel type that is assigned for the CHPID is not compatible with the IOCP type 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 PCHID. In the example, the IOCP type is OSD but the PCHID 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 (Figure 8-64 on page 471).

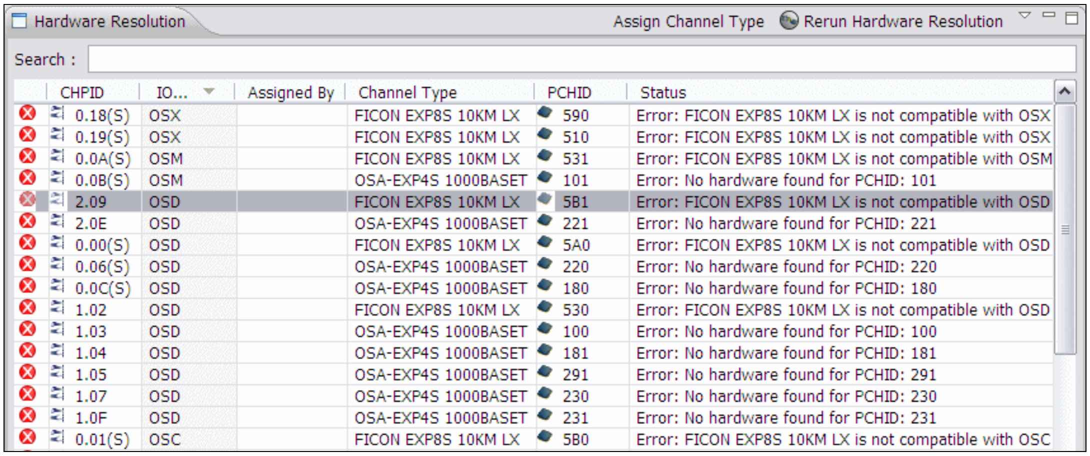

Figure 8-64 CMT - Channel_type is not compatible with IOCP_type

Select the channel type OSD. The Status is `Error: FICON EXP8S is not compatible with OSD`. Right-click in the row and select **Reset Incompatible (Hardware I/O) Entries** to remove the PCHID values for only those rows (Figure 8-65).

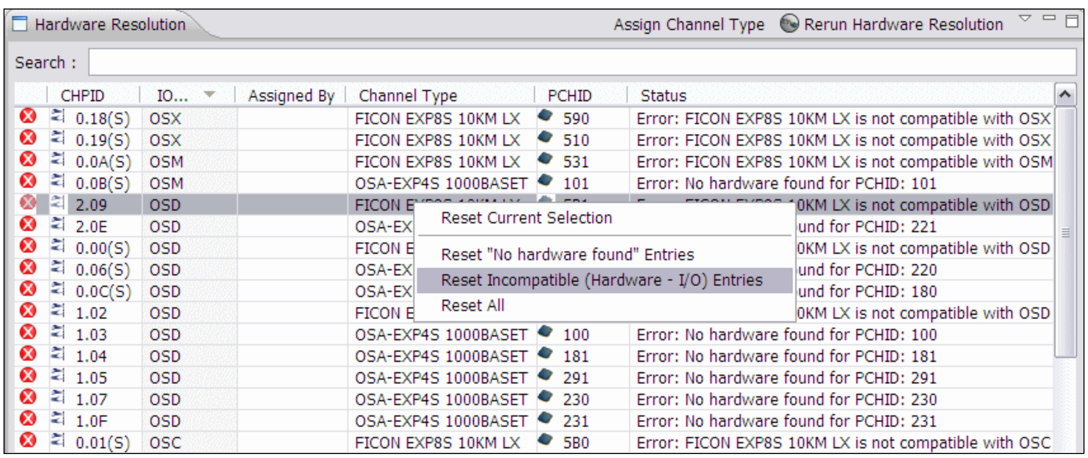

Figure 8-65 CMT - Channel_Type is not compatible with IOCP_type OSD

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 PCHIDs information (Figure 8-66).

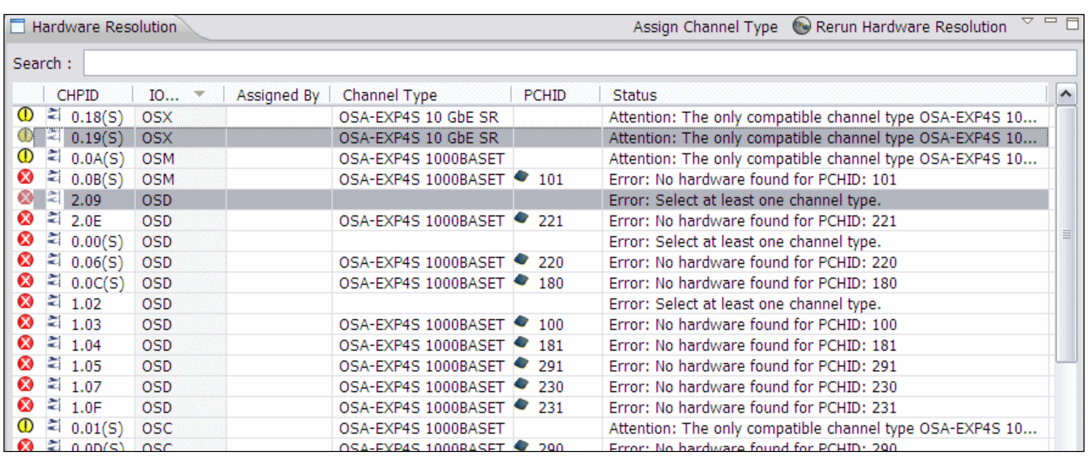

Figure 8-66 CMT - Results for reset of incompatible

The CHPID Mapping Tool now displays messages about any CHPID types that were imported from the IOCP input file (IODF) into the CMT that do not have any associated hardware support in the CFReport file (Figure 8-67). Click **OK**. The same figure also shows the Adapter Type Summary details.

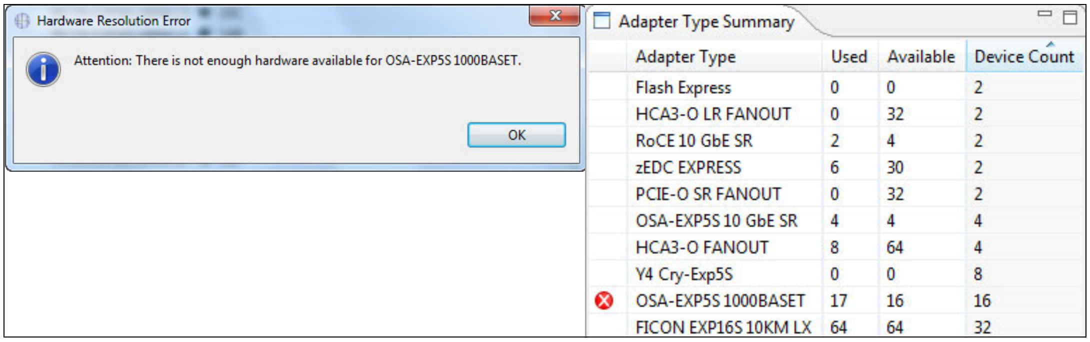

Figure 8-67 CMT - Required Hardware unavailable

Excessive numbers of *OSC* CHPID types are in the example IODF to show how the CHPID Mapping Tool handles this condition. For more information, see 8.2.3, “Overdefining channel paths on an XMP processor” on page 434.

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

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

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

To continue with this example, complete the following steps:

1. Return to the IODF and change the PCHID values for the OSD CHPIDs (or any other CHPIDs that have no 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 panel and selecting **Import IOCP File**.

Tip: If you look at the IOCP statements file now, the OSD CHPIDs have been omitted from the file, but are still defined in the IODF.

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

Reset “Error: No hardware found” Entries

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` (Figure 8-68).

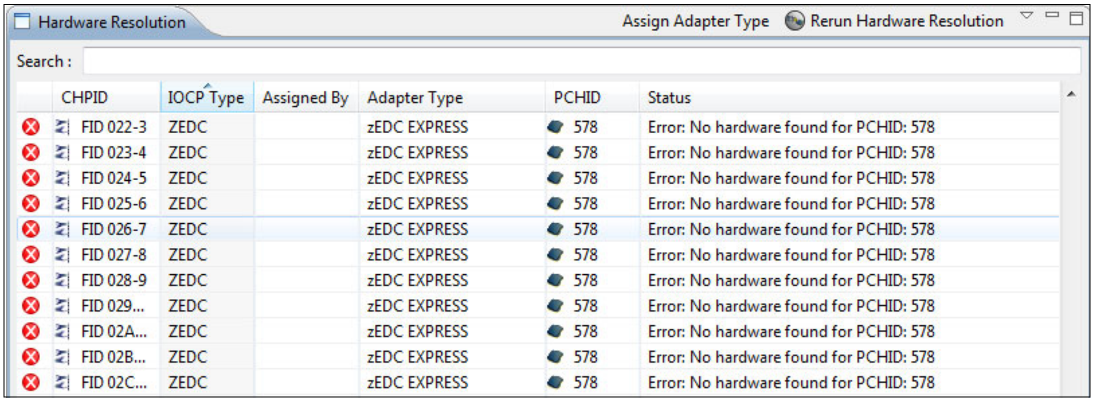

Figure 8-68 CMT - Error: No Hardware found

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 PCHID values for those rows (Figure 8-69).

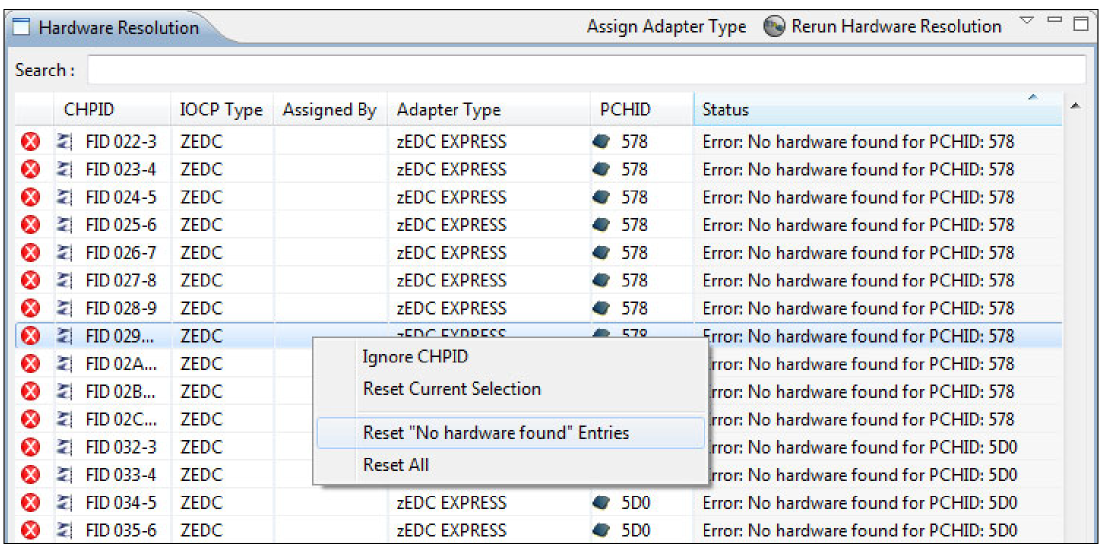

Figure 8-69 CMT - Resetting No Hardware found entries

The tool replaces the X with an *Attention* icon, changes the status message, and removes the PCHID information (Figure 8-70).

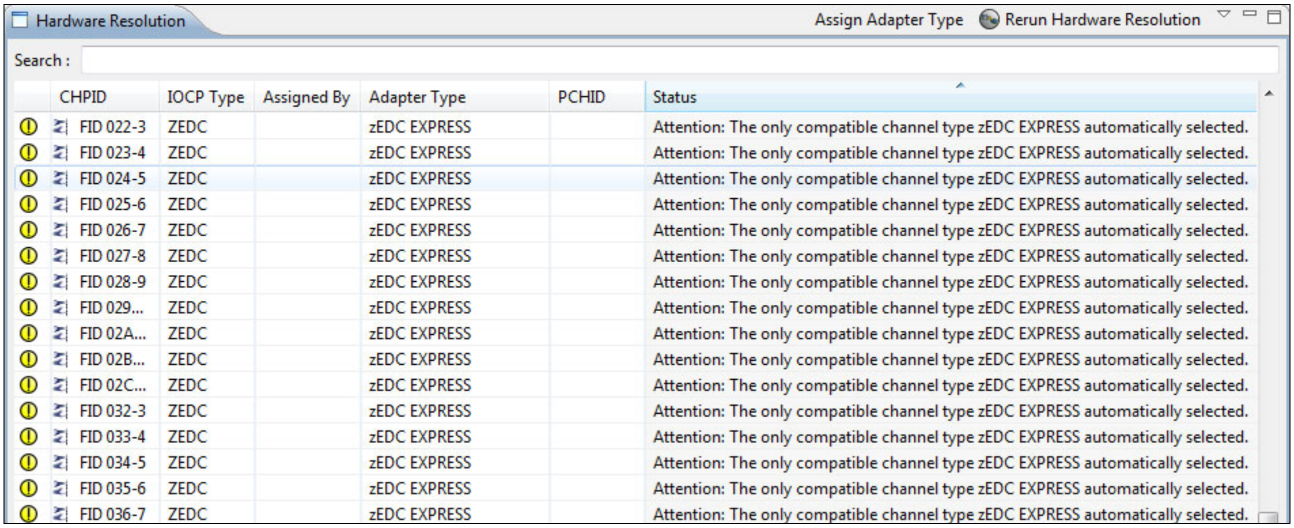

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
FID 022-3	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 023-4	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 024-5	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 025-6	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 026-7	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 027-8	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 028-9	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 029...	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 02A...	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 02B...	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 02C...	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 032-3	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 033-4	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 034-5	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 035-6	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.
FID 036-7	ZEDC		zEDC EXPRESS		Attention: The only compatible channel type zEDC EXPRESS automatically selected.

Figure 8-70 CMT - Results of resetting No hardware found

Reset “Select at least one adapter type”

The adapter type is not assigned to the current row. Assign an adapter type to the IOCP type:

1. Click the **Adapter Type** column in the target row. The tool displays an arrow in the Channel Type column of the target row (Figure 8-71).

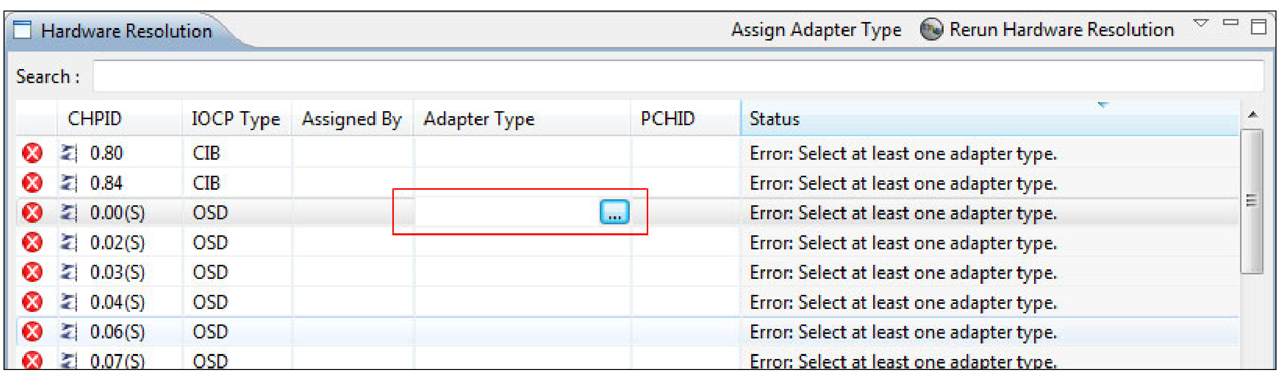

CHPID	IOCP Type	Assigned By	Adapter Type	PCHID	Status
0.80	CIB				Error: Select at least one adapter type.
0.84	CIB				Error: Select at least one adapter type.
0.00(S)	OSD		...		Error: Select at least one adapter type.
0.02(S)	OSD				Error: Select at least one adapter type.
0.03(S)	OSD				Error: Select at least one adapter type.
0.04(S)	OSD				Error: Select at least one adapter type.
0.06(S)	OSD				Error: Select at least one adapter type.
0.07(S)	OSD				Error: Select at least one adapter type.

Figure 8-71 CMT - Selecting at least one adapter type

2. Click the ellipsis (...) box.

3. The tool shows a list of available and compatible card types for the CHPID (Figure 8-72).

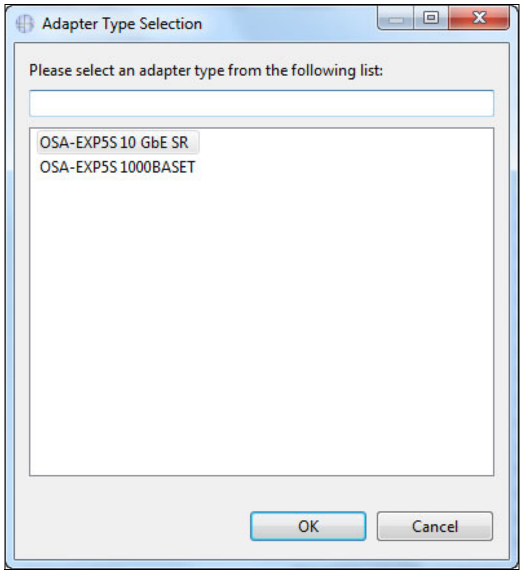

Figure 8-72 CMT - Adapter Type Selection

4. Select an adapter type and click **OK**.
5. In the Adapter Type Summary tab, observe that the “Used” and “Available” totals change.

Reset “Required hardware for type *IOCP_type* not available”

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

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

You must change IOCP file or obtain more hardware.

Reset “PCHID_1 moved to new channel ID: PCHID_2”

When moving from old hardware to new hardware, for example during a miscellaneous equipment specification (MES), the PCHID value assigned to a feature may change. This message indicates that the IOCP file contains a PCHID value for the old machine that is being removed. The PCHID value is changed from the old machine to the PCHID value for the new machine. For example, PCHID_1 is the first PCHID value representing the old hardware (for example, 1B0) and PCHID_2 is the new value representing the new hardware (for example, 533). In essence, the feature is present in both the old and new hardware, but its location (PCHID) has 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** button becomes available.

8.5.5 Manual mapping to resolve CIB CHPIDs

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

To resolve the CIB CHPIDs, assign the available CHPIDs. Click **Manual Mapping** (Figure 8-73).

Figure 8-73 CMT - Manual Mapping

Ensure that the tool is set to display Manual Mapping in **Hardware -> I/O** format (Figure 8-74).

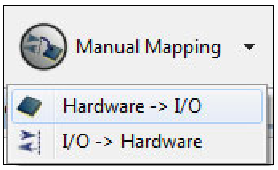

Figure 8-74 CMT - Manual Mapping of Hardware -> I/O

Click every row that has type HCA3-O in the Channel Type column. The tool displays all available CHPIDs with IOCP type (Figure 8-75).

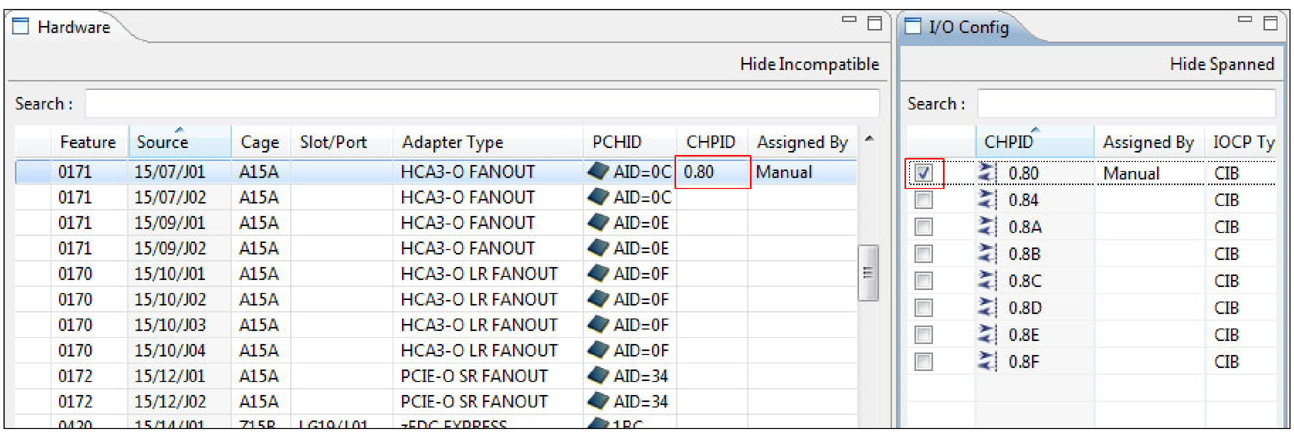

Figure 8-75 CMT - Adapter Type of HCA3 and associated CHPID assigned

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 a HCA3-O adapter type, you see the `Multiple -->` value (Figure 8-76) inserted in the CHPID and Assigned By columns.

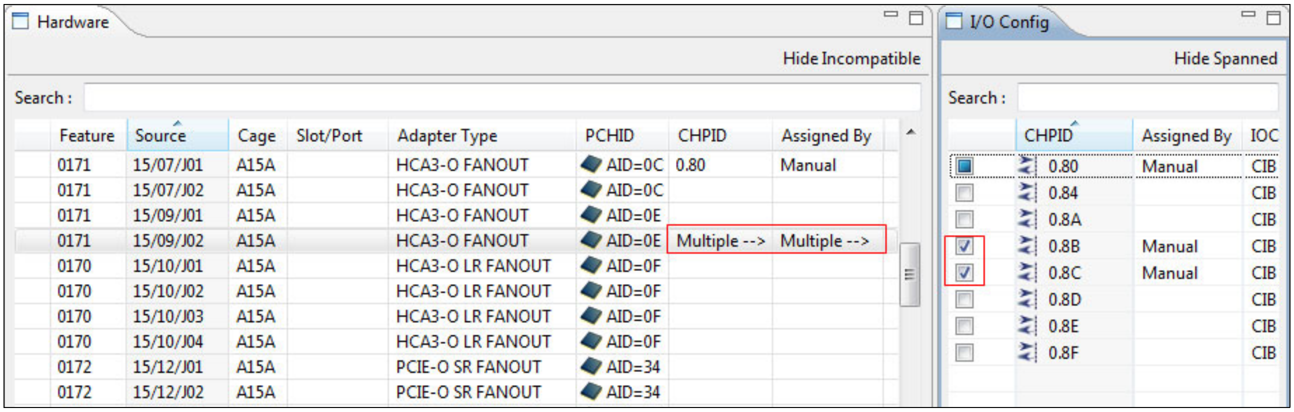

Figure 8-76 CMT - Adapter Type of HCA3 and associated multiple CHPID assigned

The **Automatic Mapping** button becomes available after you assign all the CHPIDs of IOCP type CIB.

8.5.6 Processing the CU Priority and Automatic Mapping

If you are importing an IOCP statements file from a 2817 or 2827 that had CU Priority values defined, review the CU Priority values beforehand. The CHPID Mapping Tool can then perform the availability functions appropriately for a 2964.

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 these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices (Figure 8-77 on page 478). For the example, select the following two options and then click **OK**:
   - **Reset CHPIDs assigned by Automatic Mapping**
   - **Reset CHPIDs assigned by IOCP**

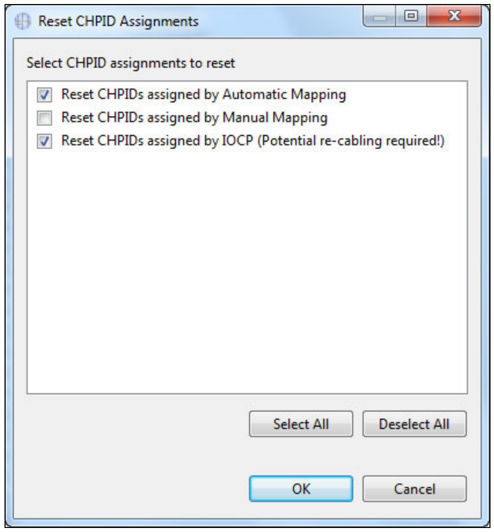

Figure 8-77 CMT - Reset CHPID Assignments

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 again to confirm the reset (Figure 8-78).

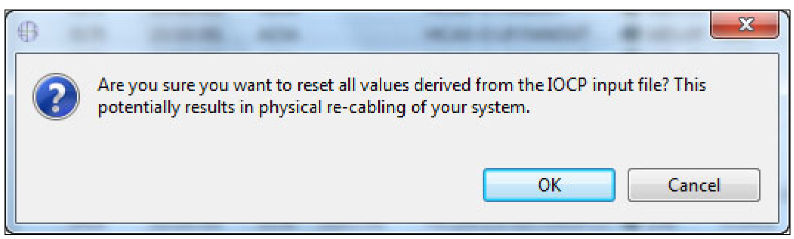

Figure 8-78 CMT - Reset CHPID assignments warning message

4. The 2964 has availability rules that differ from 2817 and 2827, so remove all PCHID 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 (Figure 8-79). Click **OK**.

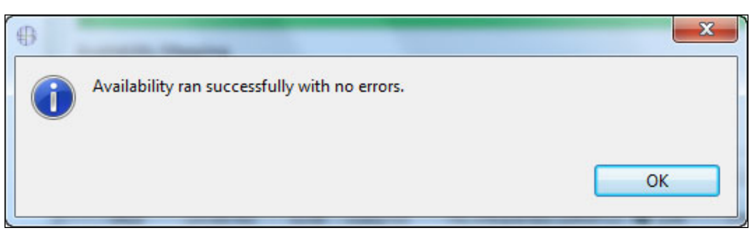

Figure 8-79 CMT - Availability ran successfully with no errors message

7. Click **OK** (Figure 8-80).

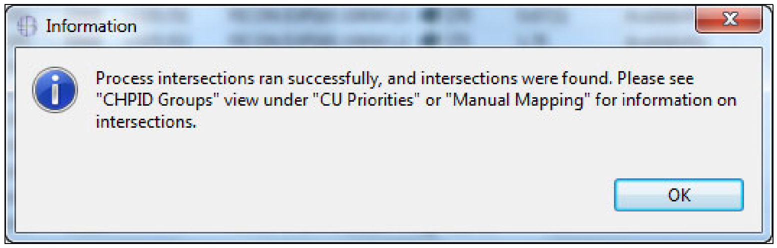

Figure 8-80 CMT - Process Intersections run successfully message

The possible intersects are as follows:

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 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 if they are acceptable (Figure 8-81). The example returned the "B" intersect.

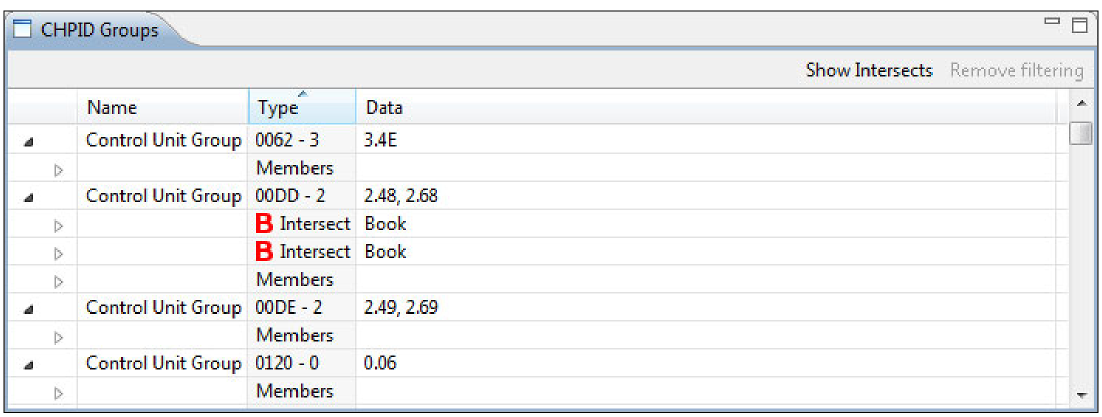

Figure 8-81 CMT - B Intersect examples

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.

Check and set values for items such as OSA-ICC CHPIDs and FCTC CHPIDs to ensure that the CHPID Mapping Tool allocates these CHPIDs with high PCHID availability.

1. Click **CU Priorities**. By default, this pane is the center on the top.
2. In the CU Priorities pane, search in the CU Number column for the control units that you want to set a priority for.

3. Type 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. In the example, set the OSC type CU Numbers to priority 333 (Figure 8-82).

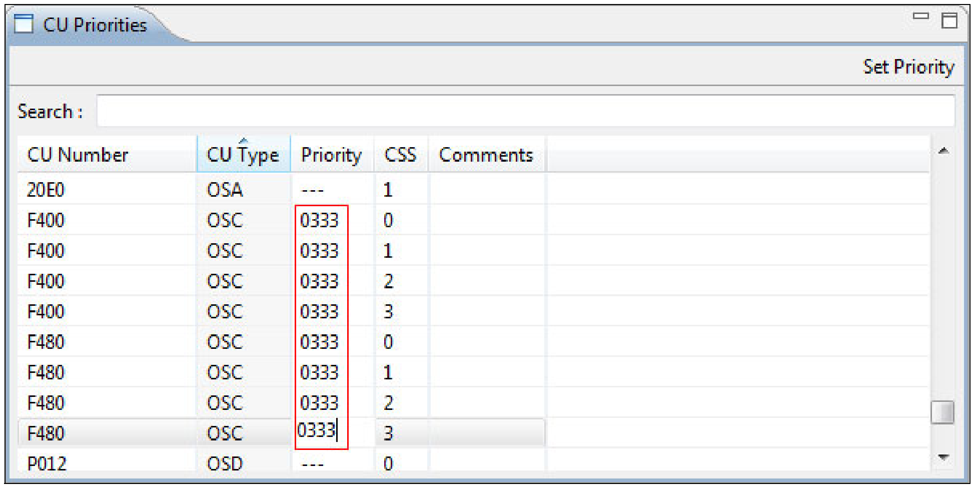

Figure 8-82 CMT - Set CU Priorities

If coupling links are used by a CF image, group those links.

Group each set of CHPIDs going to a different CPC with a common priority. For example, suppose the CF image has four links (CHPIDs 40, 41, 42, and 43) and that 40 and 41 go to one CPC, and 42 and 43 go to a different CPC. In this case, give CHPIDs 40 and 41 one priority and CHPIDs 42 and 43 a different priority. The concept is the same regardless of the number of connecting CPCs or the number of links to each CPC.

Now perform the second availability function by completing these steps:

1. Click **Automatic Mapping**.
2. The Reset CHPID Assignments window opens with Reset choices. Click **Reset CHPIDs assigned by Automatic Mapping**.
3. Click **OK**.

   The Hardware Resolution pane shows that the CHPID and Assigned By columns are no longer blank (Figure 8-83). The CMT assigned CHPIDs to PCHIDs and placed the Availability value in the Assigned By column, indicating that the CHPID values were assigned based on availability.

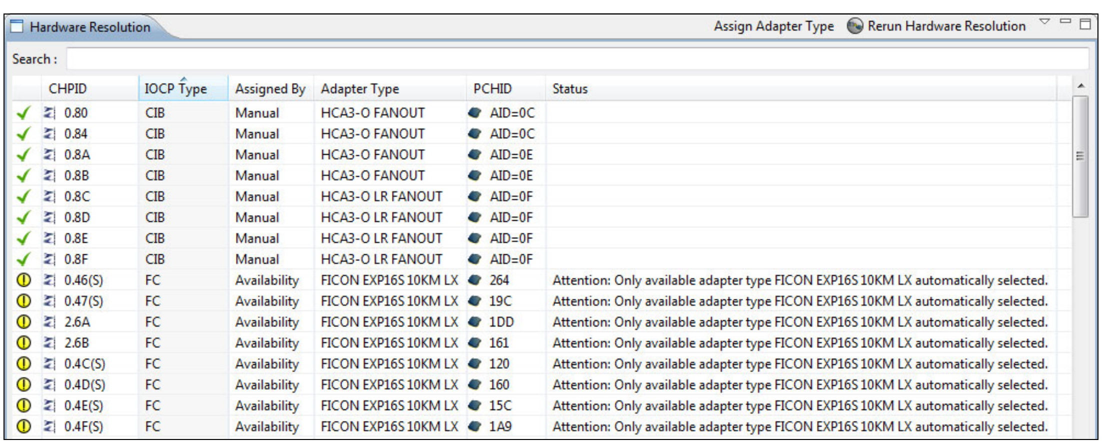

Figure 8-83 CMT - CHPIDs assigned

The possible Assigned By column values are as follows:

Manual You made the assignment by using manual mapping.

Automatic You made the assignment by using automatic mapping.

IOCP The IOCP source made the assignment.

Config File The CHPID Mapping Tool forced an assignment because of configuration file requirements.

> **More information:** See the *CHPID Mapping Tool User's Guide*, GC28-6947-00.

4. You can now display the results of the channel mapping. You can also sort the report in various ways. For example, to see how the CHPID Mapping Tool ranked the control units, select the **CU Priorities** pane and click the **Priority** column (Figure 8-84).

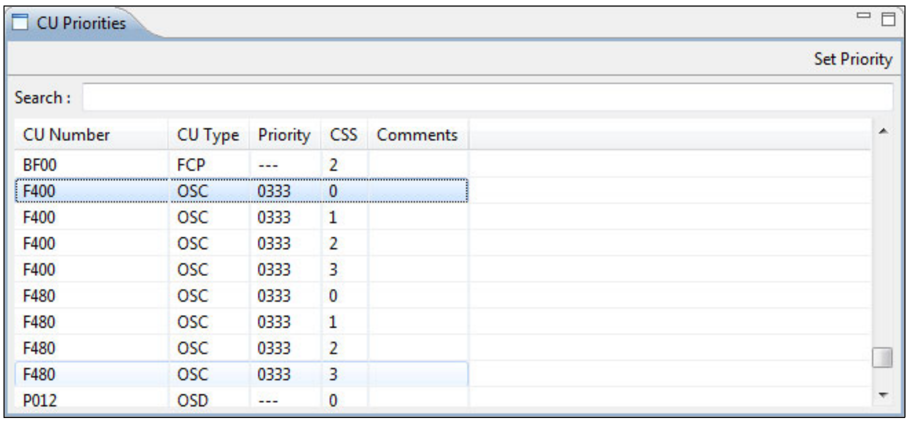

Figure 8-84 CMT - CU Priorities showing assigned priorities

The example illustrates how CU Priority values are represented in the IOCP file.

> **Tip:** The control unit priorities are stored in the IOCP output file created by the CMT that gets migrated back into HCD. HCD maintains these priorities and outputs them when it creates another IOCP deck. They are in the form of commented lines at the end of the IOCP deck, as shown here:

```
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 2341.1=0000,2361.1=0000,7800.0=0000,7800.1=0000,7800.2=0000
*CMT* 7800.3=0000,F400.0=0333,F400.1=0333,F400.2=0333,F400.3=0333
*CMT* F480.0=0333,F480.1=0333,F480.2=0333,F480.3=0333,PF030.0=0001
*CMT* PF031.0=0001,PF020.0=0001,PF032.0=0001,PF021.0=0001,PF010.0=0001
*CMT* PF022.0=0001,PF000.0=0001
```

8.5.7 CHPIDs not connected to control units

In the CU Priorities window, click in the **CU Number** column (Figure 8-85). The CHPID Mapping Tool shows, at the end of the list, all CHPIDs 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.

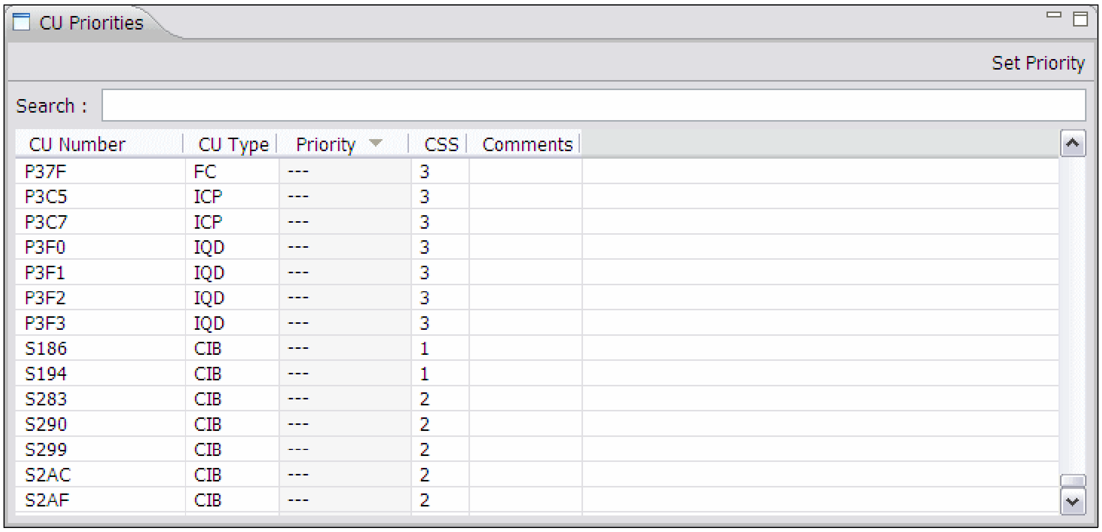

CU Number	CU Type	Priority	CSS	Comments
P37F	FC	---	3	
P3C5	ICP	---	3	
P3C7	ICP	---	3	
P3F0	IQD	---	3	
P3F1	IQD	---	3	
P3F2	IQD	---	3	
P3F3	IQD	---	3	
S186	CIB	---	1	
S194	CIB	---	1	
S283	CIB	---	2	
S290	CIB	---	2	
S299	CIB	---	2	
S2AC	CIB	---	2	
S2AF	CIB	---	2	

Figure 8-85 CMT - 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 need to 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 there are extra CHPIDs for anticipated new control units, you might want to group these CHPIDs with a common priority. Having a common priority allows the availability mapping function to pick PCHIDs that can afford your new control unit availability.

8.5.8 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**. Figure 8-86 shows the options to create a Preview Report or Save Report.

Figure 8-86 CMT - 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 (Figure 8-87).

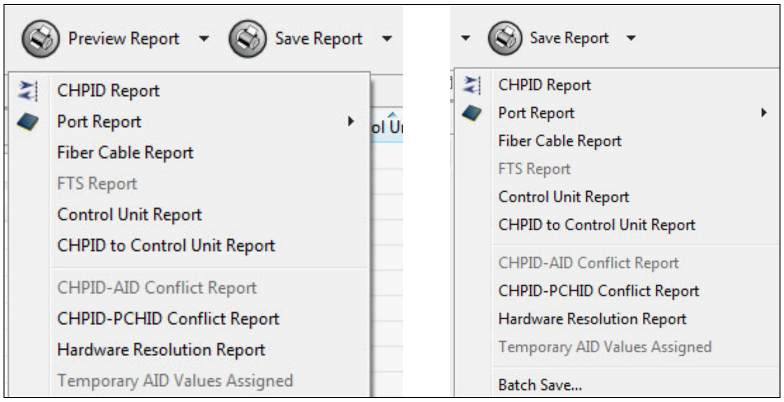

Figure 8-87 CMT - Preview Report and Save Report menus

For simplicity, only three reports are described in this example: the CHPID Report; the Port Report, sorted by location; and the 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 PCHID or cage/slot/port information before system delivery.

CHPID Report

To create the CHPID report, complete the following steps:

1. Click **Preview Report** → **CHPID Report** (Figure 8-88).

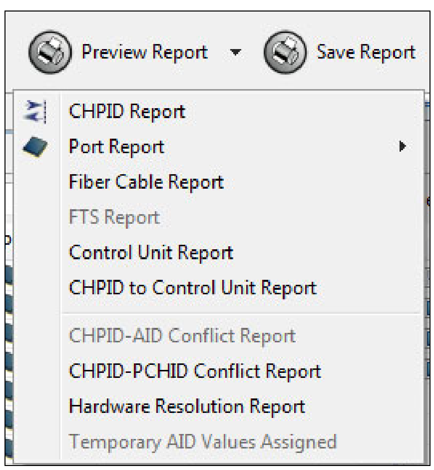

Figure 8-88 CMT - Preview report - CHPID Report

The CHPID Mapping Tool displays the CHPID Report in a Report tab within the CMT (Figure 8-89).

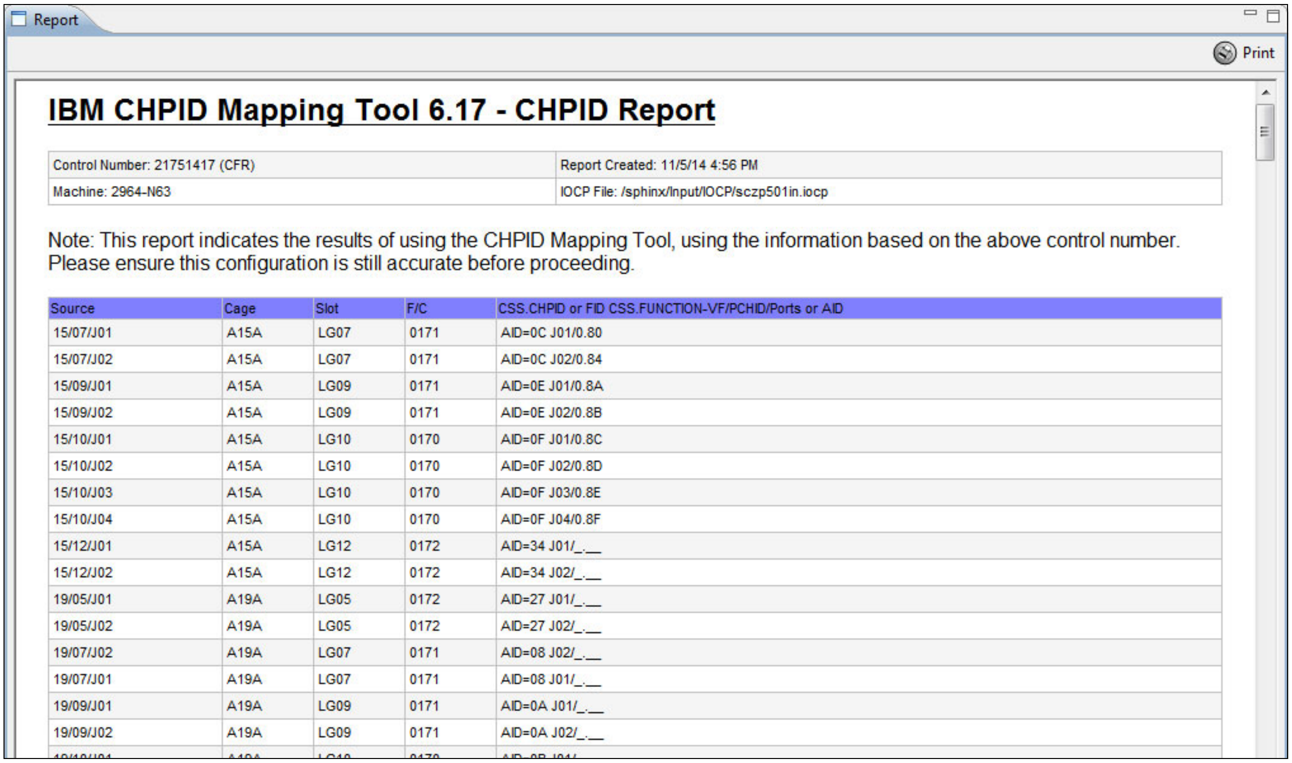
Report

Print

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 4:56 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__

Figure 8-89 CMT - 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 (Figure 8-90). 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**.

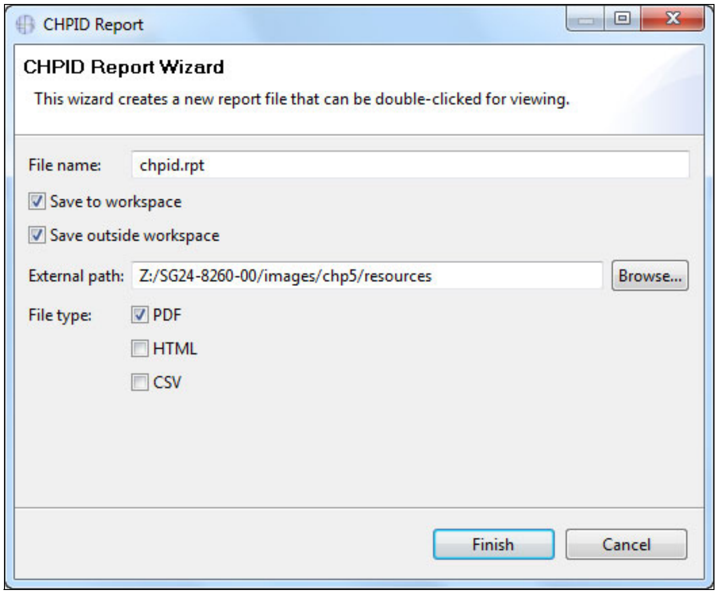

Figure 8-90   CMT - Save CHPID Report

The CHPID Report is created by the CHPID Mapping Tool (Figure 8-91).

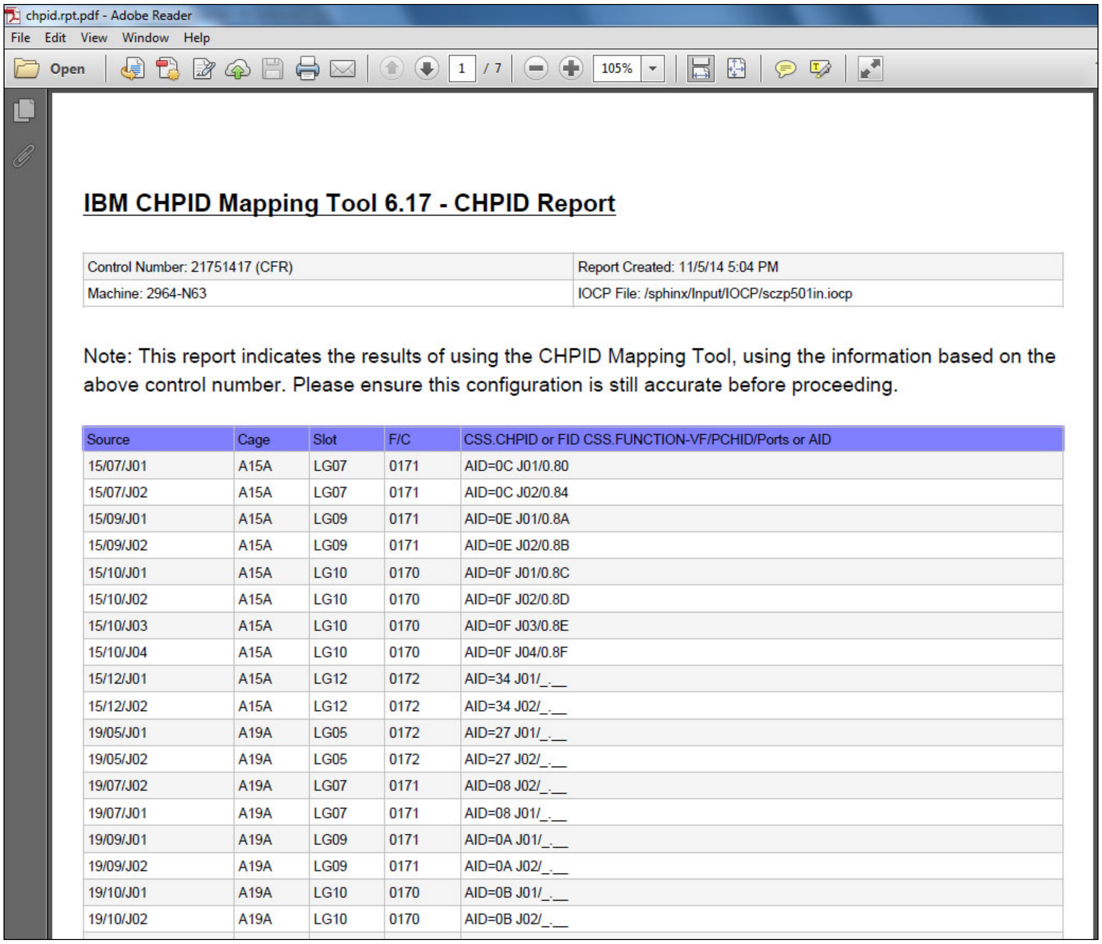

IBM CHPID Mapping Tool 6.17 - CHPID Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:04 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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.

Source	Cage	Slot	F/C	CSS.CHPID or FID CSS.FUNCTION-VF/PCHID/Ports or AID
15/07/J01	A15A	LG07	0171	AID=0C J01/0.80
15/07/J02	A15A	LG07	0171	AID=0C J02/0.84
15/09/J01	A15A	LG09	0171	AID=0E J01/0.8A
15/09/J02	A15A	LG09	0171	AID=0E J02/0.8B
15/10/J01	A15A	LG10	0170	AID=0F J01/0.8C
15/10/J02	A15A	LG10	0170	AID=0F J02/0.8D
15/10/J03	A15A	LG10	0170	AID=0F J03/0.8E
15/10/J04	A15A	LG10	0170	AID=0F J04/0.8F
15/12/J01	A15A	LG12	0172	AID=34 J01/_.__
15/12/J02	A15A	LG12	0172	AID=34 J02/_.__
19/05/J01	A19A	LG05	0172	AID=27 J01/_.__
19/05/J02	A19A	LG05	0172	AID=27 J02/_.__
19/07/J02	A19A	LG07	0171	AID=08 J02/_.__
19/07/J01	A19A	LG07	0171	AID=08 J01/_.__
19/09/J01	A19A	LG09	0171	AID=0A J01/_.__
19/09/J02	A19A	LG09	0171	AID=0A J02/_.__
19/10/J01	A19A	LG10	0170	AID=0B J01/_.__
19/10/J02	A19A	LG10	0170	AID=0B J02/_.__

Figure 8-91 CMT - CHPID Report example in PDF format

At the end of this CHPID Report is a list of CHPIDs with modified PCHID/AID assignments (Figure 8-92). This report is valuable for moving cables.

List of CHPIDs having modified PCHID/AID assignments

Note: For CHPIDs that had PCHID/AID assignments in the IOCP file that was loaded for this session of the Mapping Tool.

CHPIDs or FUNCTIONs	Previous PCHID/AID-Port	PCHID/AID-Port	Current Location	F/C
0.80	09-1	0C-1	A15ALG07J.01	0171
0.84	09-1	0C-02	A15ALG07J.02	0171
0.8A	09-1	0E-1	A15ALG09J.01	0171
0.8B	09-1	0E-02	A15ALG09J.02	0171
0.8C	09-1	0F-01	A15ALG10J.01	0170
0.8D	09-1	0F-02	A15ALG10J.02	0170
0.8E	09-1	0F-03	A15ALG10J.03	0170
0.8F	09-1	0F-04	A15ALG10J.04	0170
FID 000	544	208	Z08BLG03J.01	0411
FID 010	5EC	140	Z22BLG20J.01	0411
FID 020-1	578	1BC	Z15BLG19J.01	0420
FID 021-2	578	27C	Z08BLG38J.01	0420
FID 022-3	578	Not Assigned		
FID 023-4	578	Not Assigned		

Figure 8-92 CMT - List of CHPIDs that have modified PCHID/AID assignments

Port Report, sorted by location

To create the Port Report, 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 (Figure 8-93).

IBM CHPID Mapping Tool 6.17 - CHPID to Port Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:12 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
A15A	07/LG	AID=0C / J.01	15/07	HCA3-O FANOUT	0.80	Manual
A15A	07/LG	AID=0C / J.02	15/07	HCA3-O FANOUT	0.84	Manual
A15A	09/LG	AID=0E / J.01	15/09	HCA3-O FANOUT	0.8A	Manual
A15A	09/LG	AID=0E / J.02	15/09	HCA3-O FANOUT	0.8B	Manual
A15A	10/LG	AID=0F / J.01	15/10	HCA3-O LR FANOUT	0.8C	Manual
A15A	10/LG	AID=0F / J.02	15/10	HCA3-O LR FANOUT	0.8D	Manual
A15A	10/LG	AID=0F / J.03	15/10	HCA3-O LR FANOUT	0.8E	Manual
A15A	10/LG	AID=0F / J.04	15/10	HCA3-O LR FANOUT	0.8F	Manual
A15A	12/LG	AID=34 / J.01	15/12	PCIE-O SR FANOUT		
A15A	12/LG	AID=34 / J.02	15/12	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.01	19/05	PCIE-O SR FANOUT		
A19A	05/LG	AID=27 / J.02	19/05	PCIE-O SR FANOUT		
A19A	07/LG	AID=08 / J.01	19/07	HCA3-O FANOUT		
A19A	07/LG	AID=08 / J.02	19/07	HCA3-O FANOUT		
A19A	09/LG	AID=0A / J.01	19/09	HCA3-O FANOUT		

Figure 8-93 CMT - CHPID to Port Report, sorted by location

CHPID to CU Report

This report is created in a 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 (Figure 8-94).

IBM CHPID Mapping Tool 6.17 - CHPID to CU Report

Control Number: 21751417 (CFR)	Report Created: 11/5/14 5:15 PM
Machine: 2964-N63	IOCP File: /sphinx/Input/IOCP/sczp501in.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
0	00	OSD	19/12/J01	Z08B LG07 J00J01	214	2040	OSA	---
0	01	OSC	19/03/J01	Z15B LG21 J00J01	1C4	F400	OSC	0333
0	02	OSD	15/03/J01	Z15B LG35 J00J01	1F0	2080	OSA	---
0	03	OSD	15/15/J01	Z22B LG17 J00J01	134	20A0	OSA	---
0	04	OSD	15/02/J01	Z22B LG36 J00J01	174	20C0	OSA	---
0	06	OSD	19/02/J01	Z22B LG21 J00J01	144	0120	OSA	---
0	07	OSD	15/02/J01	Z22B LG37 J00J01	178	0140	OSA	---
0	0C	OSD	19/03/J01	Z15B LG22 J00J01	1C8	2060	OSA	---
0	0D	OSC	15/15/J01	Z22B LG18 J00J01	138	F480	OSC	0333
0	12	OSD	15/05/J01	Z08B LG36 J00J01	274	P012	OSD	---
0	13	OSD	19/14/J01	Z15B LG06 J00J01	190	P013	OSD	---
0	40	FC	15/14/J01	Z15B LG12 D1	1A4	1000	2107	---
						1200	2107	---
						1400	2107	---

Figure 8-94 CMT - CHPID to CU Report

8.5.9 Creating an updated IOCP

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

> **Using HCM:** You might prefer to use HCM to transfer the updated IOCP statements file back to the host. However, first run the next step in the CHPID Mapping Tool to create the updated IOCP file.

To create the IOCP, complete the following steps:

1. Select **File** → **Export IOCP input file** (Figure 8-95 on page 489).

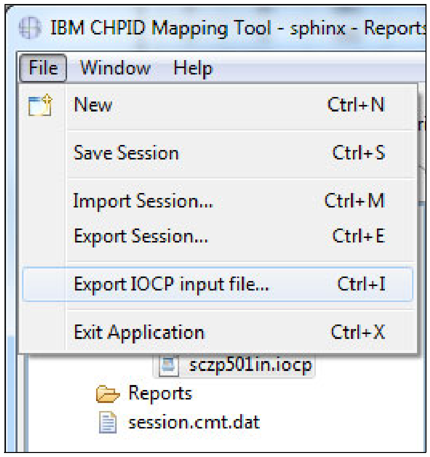

Figure 8-95 CMT - Export IOCP input file

2. Enter the Export Path and IOCP Name for the IOCP output file and click **Finish** (Figure 8-96).

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

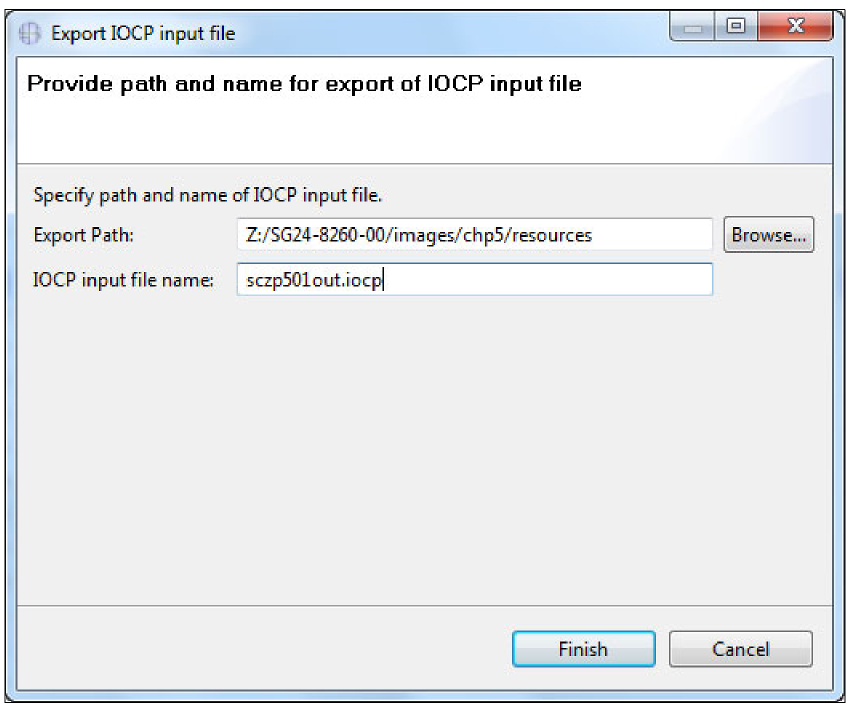

Figure 8-96 CMT - Export IOCP File

3. Select **File** → **Save Session** (Figure 8-97 on page 490).

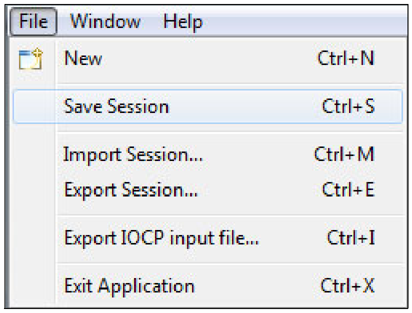

Figure 8-97   CMT - Save session

You might want to save your project before exiting the CMT application.

8.6 HCD: Updating the 2964 work IODF with PCHIDs

After you map the PCHIDs to CHPIDs by using the CHPID Mapping Tool, transfer this information back into HCD. To update the IODF with the PCHIDs, follow these following steps:

1. Upload the IOCP file that was created by the CMT (`SCZP501out.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 has left a reference to the CCN. Also, note the CU Priority values 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. They are in the form of commented lines at the end of the IOCP deck (Example 8-4).

Example 8-4   HCD - Updated IOCP statements file with CMT statements

```
*CMT* VERSION=000
*CMT* CCN=21751417(CFR from ResourceLink)
*CMT* 0061.0=0400,0062.0=0400,0B00.0=0500,0B01.0=0500,0B20.0=0520
*CMT* 1000.0=0100,1100.0=0100,1200.0=0100,1300.0=0100,1400.0=0100
*CMT* 1500.0=0100,1600.0=0100,1700.0=0100,2100.0=0320,2120.0=0320
*CMT* 4B00.0=0600,4B08.0=0600,5B00.0=0600,5B08.0=0600,9000.0=0120
*CMT* 9080.0=0120,9100.0=0120,9180.0=0120,9200.0=0120,9280.0=0120
*CMT* 9300.0=0120,9380.0=0120,9400.0=0120,9480.0=0120,9500.0=0120
*CMT* 9580.0=0120,9600.0=0120,9680.0=0120,9800.0=0120,9880.0=0120
*CMT* 9900.0=0120,9980.0=0120,9A00.0=0120,9A80.0=0120,9B00.0=0120
*CMT* 9B80.0=0120,9C00.0=0120,9C80.0=0120,9D00.0=0120,9D80.0=0120
*CMT* 9E00.0=0120,9E80.0=0120,A000.0=0100,A100.0=0100,A200.0=0100
*CMT* A300.0=0100,A400.0=0100,A500.0=0100,A600.0=0100,A700.0=0100
*CMT* C000.0=0160,C100.0=0160,C200.0=0160,C300.0=0160,C400.0=0160
*CMT* C500.0=0160,C600.0=0160,C700.0=0160,C800.0=0160,C900.0=0160
*CMT* CA00.0=0160,CB00.0=0160,CC00.0=0160,CD00.0=0160,CE00.0=0160
*CMT* CF00.0=0160,D000.0=0180,D100.0=0180,D200.0=0180,D300.0=0180
*CMT* D400.0=0180,D500.0=0180,D800.0=0180,D900.0=0180,DA00.0=0180
*CMT* DB00.0=0180,DC00.0=0180,DD00.0=0180,DE00.0=0180,DF00.0=0180
*CMT* F401.0=0300,F481.0=0300,PF030.0=0001,PF031.0=0001,PF020.0=0001
*CMT* PF032.0=0001,PF021.0=0001,PF010.0=0001,PF022.0=0001,PF011.0=0001
 *CMT* PF000.0=0001
 ******************************* Bottom of Data ***********************
```

Important: Do not edit the CMT comments that are contained in the IOCP output file manually. If priorities must be changed, use the CMT.

2. On the HCD main panel, enter the work IODF name that you used to create the IOCP input data set for the CHPID Mapping Tool (Figure 8-98). Select option **5. Migrate configuration data**.

```
                            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 . . . 'SYS6.IODFB8.WORK'                    +
```

Figure 8-98 HCD - Hardware Configuration: migrate configuration data

3. From the Migrate Configuration Data panel (Figure 8-99), 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
```

Figure 8-99 HCD - Migrate Configuration Data: migrate IOCP/OS data

4. The Migrate IOCP Data panel opens (Figure 8-100 on page 492). 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 OS configuration that is associated with the processor

IOCP only input data set This data set was specified when the `sczp501out.iocp` file was uploaded to z/OS.

Processing mode Use option **2** to save the results of the migration. (Before using option 2, try to migrate using option **1** to validate the operation.)

Migrate Options Use 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 . . . . . . . . . . . . SCZP501_  +   CSS ID . . . . . . . _  +
OS configuration ID  . . . . . . . . L06RMVS1  +

Combined IOCP/MVSCP input data set . _____________________________________
IOCP only input data set . . . . . . 'SYS6.IODFB8.IOCPOUT.SCZP501'
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)
```

Figure 8-100 HCD - 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 (Figure 8-101).

```
------------------------- 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
#                     SYS6.IODFB8.WORK.
**************************** Bottom of data *****************************
```

Figure 8-101 HCD - Migration Message List: successful message

The work IODF now contains both the CHPID definitions and the mapping to PCHIDs 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.

8.7 HCD: Building the 2964 production IODF

To use the definitions that were updated in HCD, create a production IODF from the work IODF.

Complete the following steps:

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 8-102).

```
                               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.IODFB8.WORK'                    +
```

Figure 8-102 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 8-103). 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
```

Figure 8-103 HCD - Activate or Process Configuration Data: build production IODF

3. HCD displays the Message List panel (Figure 8-104). Verify that you have only severity "W" (warning) messages and that they are normal for the configuration. Correct any other messages that should not occur and try to build the production IODF again. Continue this process until you have no messages that indicate problems.

```
------------------------------ Message List ------------------------------
  Save  Query  Help
---------------------------------------------------------------------------
                                                                Row 1 of 173
Command ===> ___________________________________________ Scroll ===> CSR

Messages are sorted by severity. Select one or more, then press Enter.
/ Sev Msg. ID  Message Text
_ W   CBDG159I Switch control unit(s) 0061 and device(s) 0061 defined,
#              but not yet connected to both a processor and an
#              operating system.
_ W   CBDG159I Switch control unit(s) 0062 and device(s) 0062 defined,
#              but not yet connected to both a processor and an
#              operating system.
_ W   CBDG159I Switch control unit(s) 0063 and device(s) 0063 defined,
#              but not yet connected to both a processor and an
#              operating system.
```

Figure 8-104 HCD - Message List: building the production IODF

4. Press PF3 to continue.
5. The Build Production I/O Definition File panel opens (Figure 8-105). Complete the Production IODF name and Volume serial number fields, and then press Enter.

```
------------- Build Production I/O Definition File ----------


Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODFB8.WORK'

Production IODF name . 'SYS6.IODFB8'__________________
Volume serial number . IODFPK  +

Continue using as current IODF:
2   1.  The work IODF in use at present
    2.  The new production IODF specified above
```

Figure 8-105 HCD - Build Production I/O Definition File: data fields to be updated

6. The Define Descriptor Fields panel opens (Figure 8-106). 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.IODFB8'

Descriptor field 1  . . . SYS6
Descriptor field 2  . . . IODFB8
```

Figure 8-106 HCD - 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.IODFB8 created.`

8.8 HCD/HMC: Loading the 2964 processor IOCDS

You now have a production IODF, named `SYS6.IODFB8`. Now the IOCDS component of the IODF must be updated on the *new* CPC that is being installed (for example, SCZP501) and activated (POR) using this IOCDS.

The final step is to perform an initial program load (IPL) of the processor using this IODF. (Describing how to perform the IPL of the new hardware is beyond the scope of this book.)

The two possible ways to load the IOCP Statements onto the 2964 Support Element IOCDS are as follows:

- HCD, using option 2.11
- The HMC/SE, by using the Stand-Alone Input/Output Configuration Program

Although both are valid methods to write the new configuration to the IOCDS, using HCD option 2.11 is the preferred method. However, your new 2964 processor and Support Element might not be connected to the system where the configuration was generated or to any system where HCD is running. In that case, you must use the Stand-Alone IOCP process.

8.8.1 Updating the IOCDS using HCD option 2.11

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

1. From the HCD main menu, select option **2. Activate or process configuration data** (Figure 8-107). Ensure that the IODF is the production one created in 8.7, “HCD: Building the 2964 production IODF” on page 493. Press Enter.

```
                          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.IODFB8'                      +
```

Figure 8-107 HCD - Hardware Configuration: activate or process configuration data

2. The Activate or Process Configuration Data panel opens (Figure 8-108).
Select option **11. Build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes**.

```
----- 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
```

Figure 8-108 HCD - Activate or Process Configuration data: build and manage System z cluster IOCDSs, IPL attributes and dynamic I/O changes

Consideration: This example assumes that you have connectivity to the new 2964 over the HMC LAN to create an IOCDS from which you power-on reset.

If the new 2964 is not accessible from the HMC LAN, copy the IOCP statements onto a USB flash memory drive. You can then import them into the 2964 HMC to run a stand-alone IOCP. Creating a file on a USB flash memory drive can be done by using the same process that is used to create an IOCP input file for the CHPID Mapping Tool.

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

3. The System z Cluster List panel opens (Figure 8-109). In the list, use a forward slash (/) to select the new 2964 to update one of its IOCDSs. Then, press Enter.

```
                           System z Cluster List                  Row 1 of 4
Command ===> ______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
_ USIBMSC.SCZP401    2827   H43     SCZP401
/ USIBMSC.SCZP501    2964   N63     SCZP501
****************************** Bottom of data *******************************
```

Figure 8-109 HCD - System z Cluster List: selecting processor for IOCDS update

4. The Actions on selected CPCs panel opens (Figure 8-110). 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)
```

Figure 8-110 HCD- Actions on selected CPCs: work with IOCDSs

5. The IOCDS List panel opens (Figure 8-111). Select the IOCDS that you want to update for the 2964 installation by typing a forward slash (/) next to it, and 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.SCZP501  IODFB7   LPAR   Alternate  No        No          No
_ A1.SCZP501  IODF40   LPAR   Alternate  No        No          No
_ A2.SCZP501  IODF42   LPAR   Alternate  No        No          No
_ A3.SCZP501           LPAR   POR        Yes       No          Yes-POR
****************************** Bottom of data *******************************
```

Figure 8-111 HCD - IOCDS List: selecting IOCDS for update

6. The Actions on selected IOCDSs panel opens (Figure 8-112). Select option **1. Update IOCDS** and press Enter.

```
--------------- Actions on selected IOCDSs ------


Select by number or action code and press Enter.

1_  1.  Update IOCDS . . . . . . . . . . . . . (u)
    2.  Switch IOCDS . . . . . . . . . . . . . (s)
    3.  Enable write protection  . . . . . . . (e)
    4.  Disable write protection . . . . . . . (w)
```

Figure 8-112 HCD - Actions on selected IOCDSs; update IOCDS

7. The Build IOCDSs panel opens (Figure 8-113). Verify that all the information is correct. Complete the Title1 field and press Enter

```
------------------------------ Build IOCDSs -------------------------------
                                                                 Row 1 of 1
Command ===> ___________________________________________ Scroll ===> CSR

Specify or revise the following values.

IODF name . . . . . . . . . : 'SYS6.IODFB8'

Title1 . IODFB8__________________________________________________________
Title2 : SYS6.IODFB8 - 2014-11-10 17:03

                           Write IOCDS in
IOCDS        Switch IOCDS  preparation of upgrade
A0.SCZP501   No            No
**************************** Bottom of data *****************************
```

Figure 8-113 HCD - Build IOCDSs: verifying IODF

8. The Job Statement Information panel opens (Figure 8-114 on page 499). Complete the job statements as required by the installation and press Enter. HCD submits the job to update the IOCDS.

```
----------------------- Job Statement Information -------------


Specify or revise the job statement information.


Job statement information
//WIOCP   JOB (ACCOUNT),'NAME',MSGCLASS=H
//*
//*
//*
//*
//*
```

Figure 8-114   HCD - Job Statement Information: option to override job statement cards

9. Verify the job output to ensure that the IOCDS was written without error and to the correct IOCDS. You receive a message similar to the following message:

```
ICP057I IOCP JOB WIOCP    SUCCESSFUL.  LEVEL A0 IOCDS REPLACED.

Sev  Msgid    Message Text
 I  CBDA674I  IOCP successfully completed for A0.SCZP501.
```

10. Now if you return to HCD option 2.11 and view the IOCDS, the SNA Address is at `USIBMSC.SCZP501` (Figure 8-115).

```
                         System z Cluster List                 Row 1 of 4
Command ===> _______________________________________________ Scroll ===> CSR

Select one or more CPCs, then press Enter.

  --------------CPC--------------   IODF
/ SNA Address        Type   Model   Processor ID
# USIBMSC.DSTCP02    2097   E64
_ USIBMSC.SCZP201    2097   E26     SCZP201
_ USIBMSC.SCZP401    2827   H43     SCZP401
s USIBMSC.SCZP501    2964   N63     SCZP501
****************************** Bottom of data *******************************
```

Figure 8-115   HCD - System z Cluster List: selecting processor for IOCDS verify

Figure 8-116 shows the updated IOCDS with Alternate status.

```
                                 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.SCZP501   IODFB8   LPAR   Alternate  No        Yes         No
_ A1.SCZP501   IODF40   LPAR   Alternate  No        No          No
_ A2.SCZP501   IODF42   LPAR   Alternate  No        No          No
_ A3.SCZP501   IODF43   LPAR   POR        Yes       No          Yes-POR
****************************** Bottom of data *******************************
```

Figure 8-116   HCD - IOCDS List: IOCDS verified

8.8.2 Updating the IOCDS using Stand-Alone Input/Output Config Program

Copy the IOCP statements that were generated by using HCD option **2.3. Build IOCP input data set** onto a USB flash memory drive and retain it.

> **Tip:** For more information, see the *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152.

To update the IOCDS, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, as opposed to a remote web browser. Select the new 2964, assuming it was defined to the Defined CPCs Work Area.
2. Perform a power-on reset using one of the Reset profiles and Starter IOCDSs provided with the processor during installation. This action creates an environment on the processor in which you can run the Stand-Alone IOCP process.
3. When the power-on reset is complete, activate one of the logical partitions with at least 128 MB of storage. Use this partition to run the I/O Configuration Program.
4. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
5. Under Systems or Members, click the system to select it (in this example, SCZP501).
6. On the Tasks tab, click **Recovery** → **Single Object Operations** → **Yes**.
7. Under Systems Management, click the system to select it (in this example, SCZP501).
8. Under Partitions, select the LPAR you want to use to run the Stand-Alone IOCP program (in this example, A0B) as shown in Figure 8-117 on page 501.

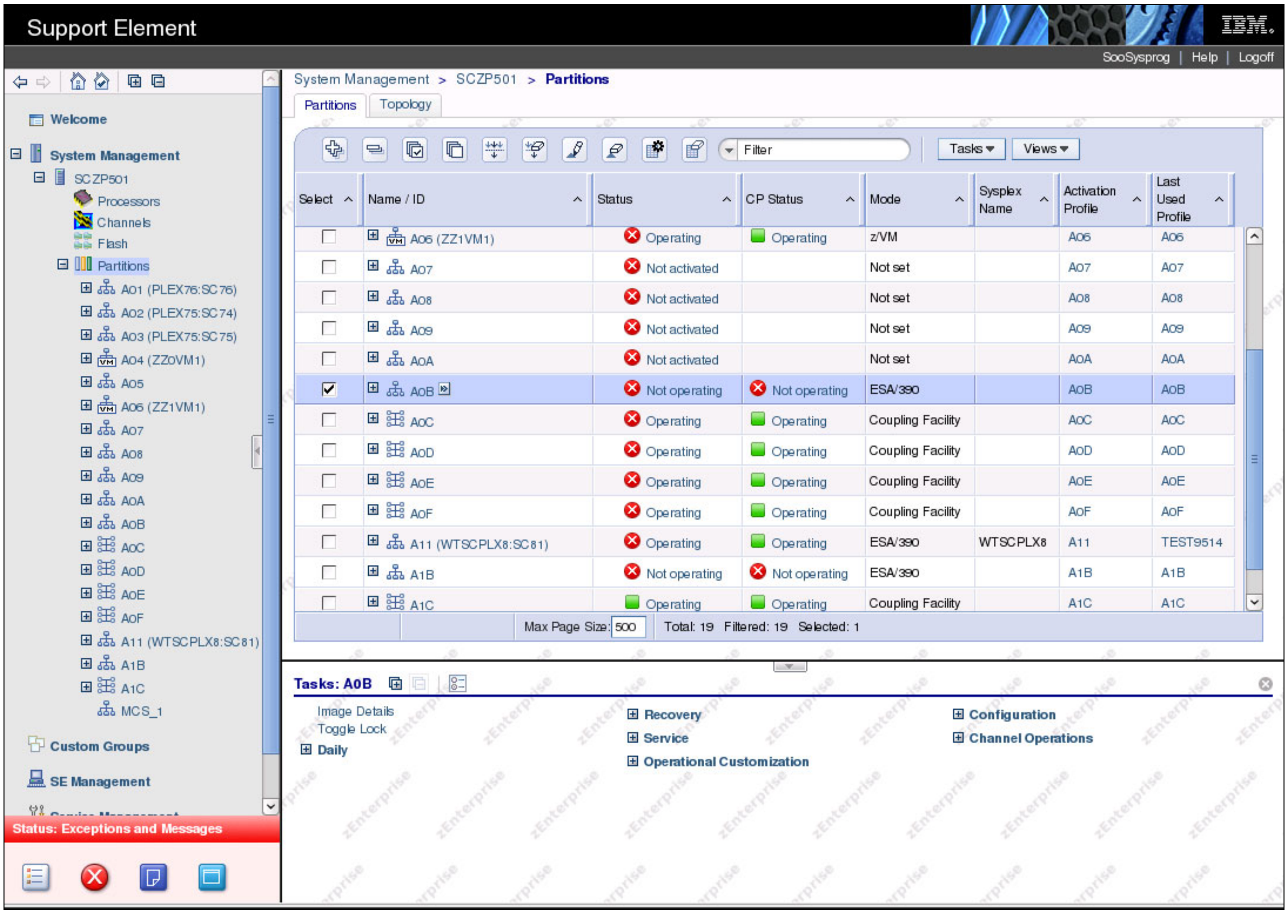

Figure 8-117   HMC - SAIOCP: partition selected for SAIOCP program load

9. On the Tasks tab, click **Configuration** → **Input/output (I/O) Configuration**.

10. Select the data set into which you want to import the IOCDS (in this example, A0) as shown in Figure 8-118.

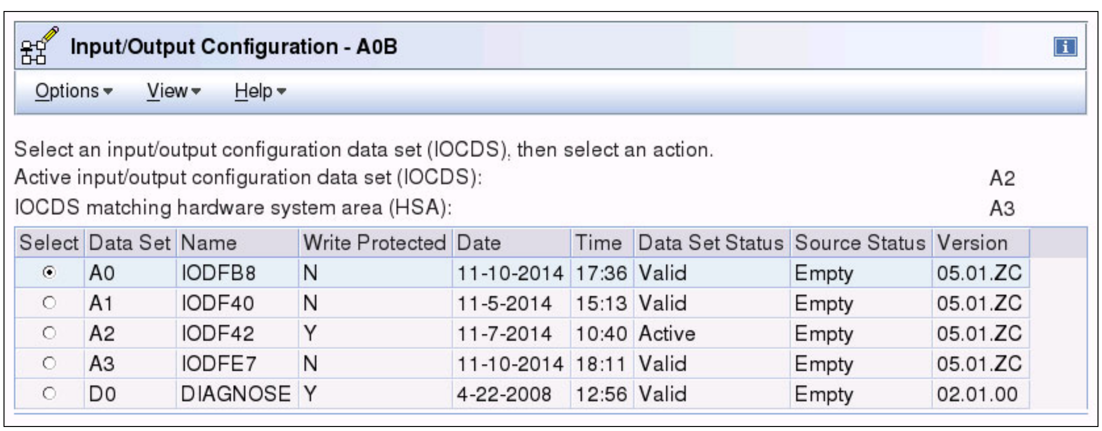

Figure 8-118   HMC - SAIOCP: IOCDS selection for import

11. Insert the USB flash memory drive that contains the IOCP text file. Wait for the drive insertion message to open (Figure 8-119).

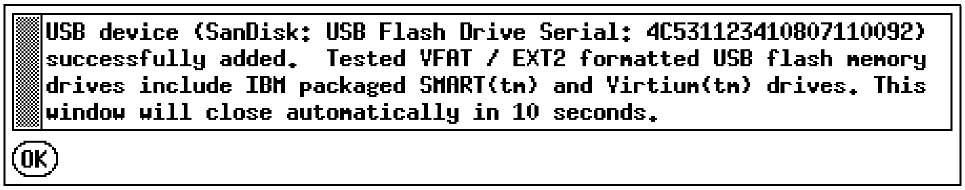

Figure 8-119 HMC - SAIOCP: USB drive insertion message

> **Tip:** Only files in the root directory of the USB Drive can be read by the HMC. Any folders and their contents will not be read.

12. Click **Options** → **Import Source File** → **Hardware Management Console USB Flash Memory Drive** (Figure 8-120).

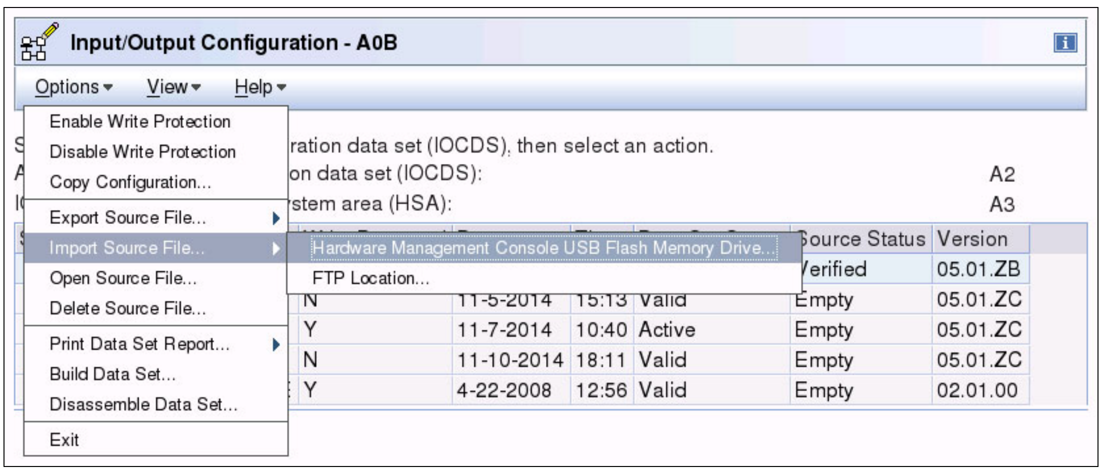

Figure 8-120 HMC - SAIOCP: import source file

13. Select the source file name and click **OK** (Figure 8-121).

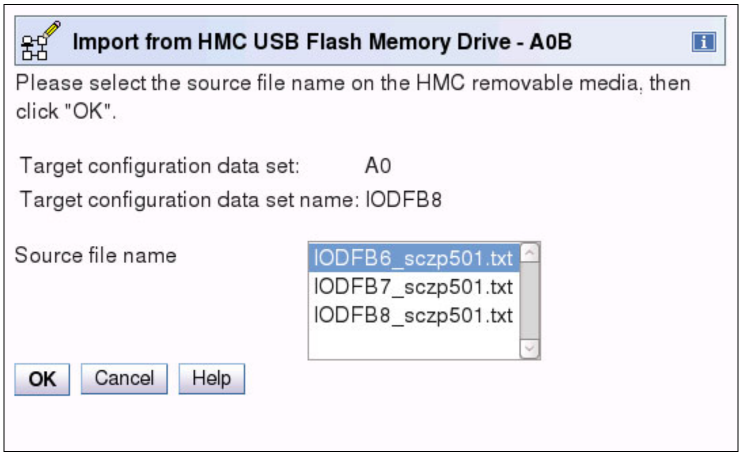

Figure 8-121 HMC - SAIOCP: select source file

14. The source file is now read from the USB drive. Click **OK** (Figure 8-122).

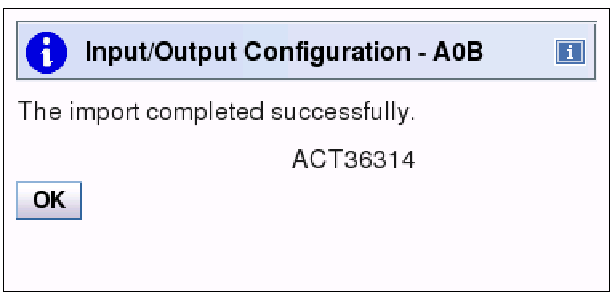

Figure 8-122 HMC - SAIOCP: source file imported

The Source Status now says `Imported` (Figure 8-123).

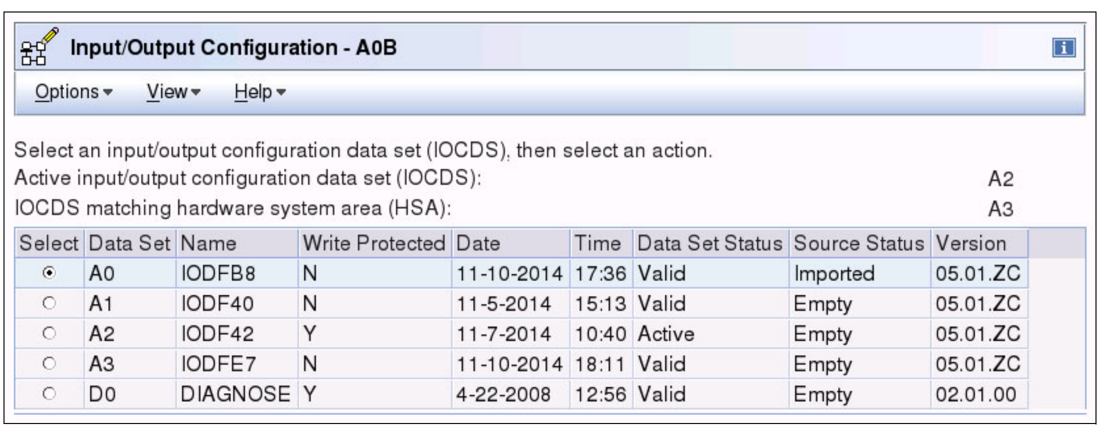

Figure 8-123 HMC - SAIOCP: IOCDS source status has changed to imported

15.Click **Options** → **Build Data Set** (Figure 8-124).

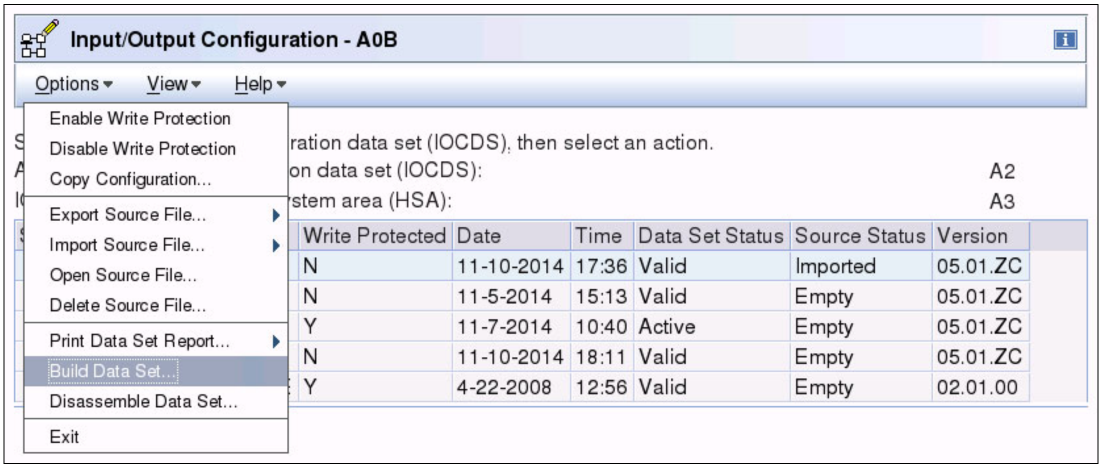

Figure 8-124 HMC - SAIOCP: building the IOCDS

16.Select the build options you want and click **OK** (Figure 8-125).

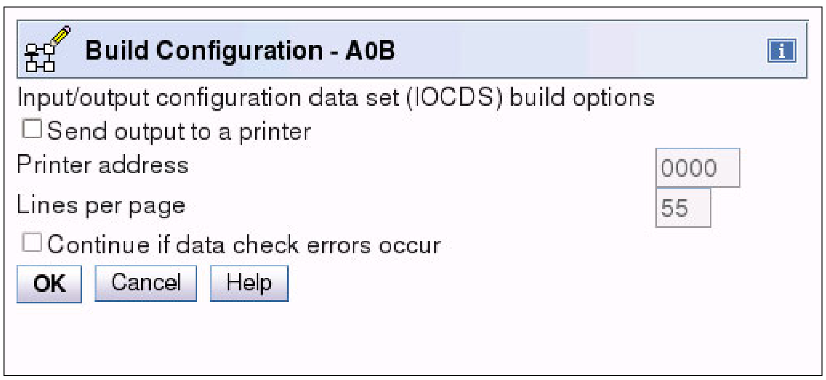

Figure 8-125 HMC - SAIOCP: build options

17. Observe the Build warning message, enter your HMC password, and click **Yes** (Figure 8-126).

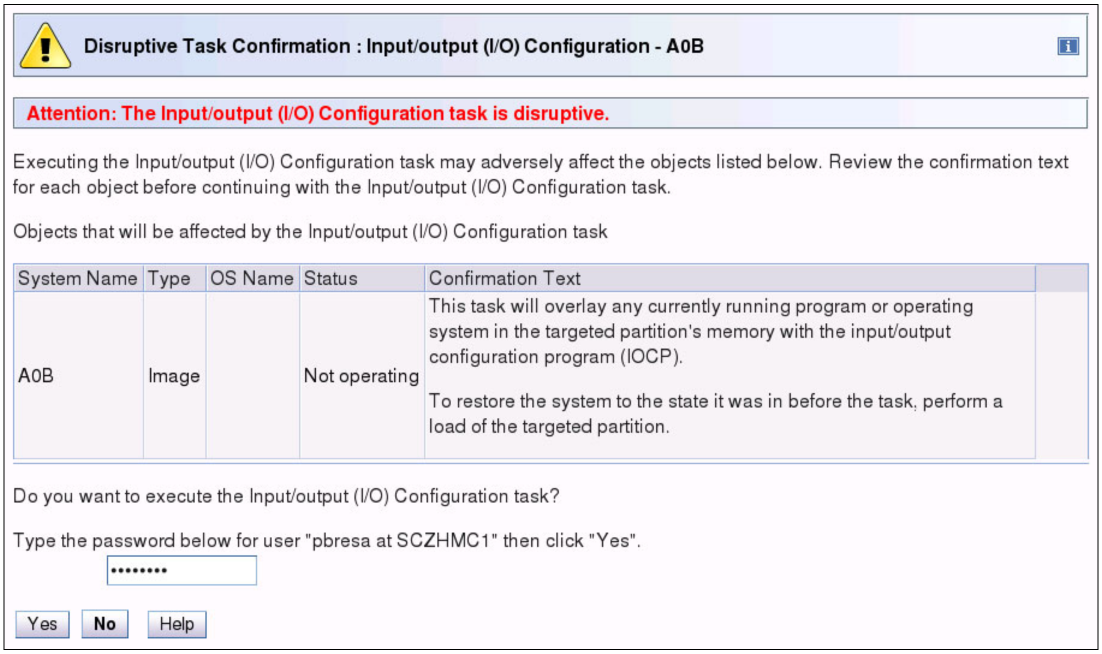

Figure 8-126 HMC - SAIOCP - build warning window

18. Status messages are displayed during the build process. After the process completes successfully, the message shown in Figure 8-127 is displayed. Click **OK**.

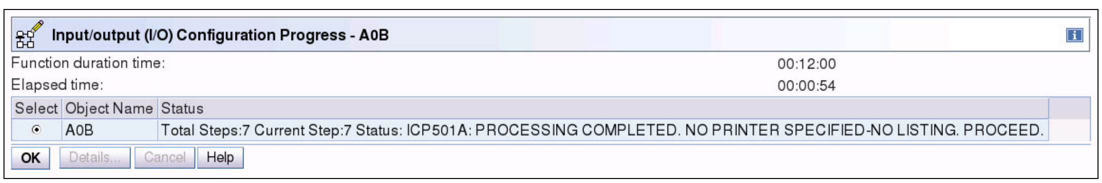

Figure 8-127 HMC - SAIOCP - build process status message

19. Observe that the Source Status now says `Verified` (Figure 8-128).

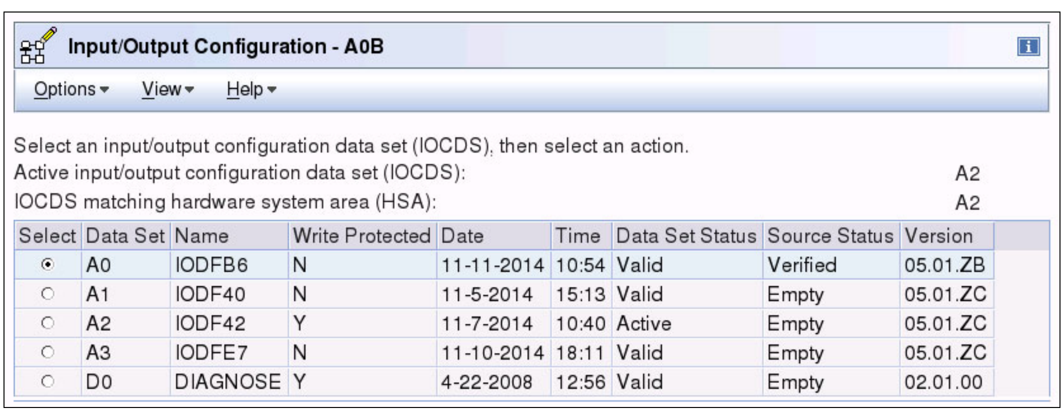

Figure 8-128 HMC - SAIOCP - IOCDS source status changed to verified

20. Click **Options** → **Exit** to end the IOCDS build process and deactivate the LPAR if it is no longer required.

This IOCDS is now ready to be selected by a Reset Profile. The 2964 can be activated (power-on reset) with the production IODF.

The USB drive can also now be removed. A drive removal message is displayed (Figure 8-129).

Figure 8-129 HMC - SAIOCP: USB drive removal message

8.9 HMC: Steps for profile activation

To activate your Reset profile using the HMC, complete the steps in this section.

8.9.1 Building the Reset Profile for Activation 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 2964 after it is upgraded and handed over from the IBM System Service Representative.

To build the profile, complete the following steps:

1. Log on using SYSPROG authority to the HMC workstation supplied with the 2964, or use a remote web browser and select the new 2964.
2. Under Systems Management or Ensemble Management, click **Systems or Members** to expand the list.
3. Under Systems or Members, click the system to select it (in this example, SCZP501).
4. On the Tasks tab, click **Operational Customization** to expand it, and select **Customize/Delete Activation Profiles** (Figure 8-130).

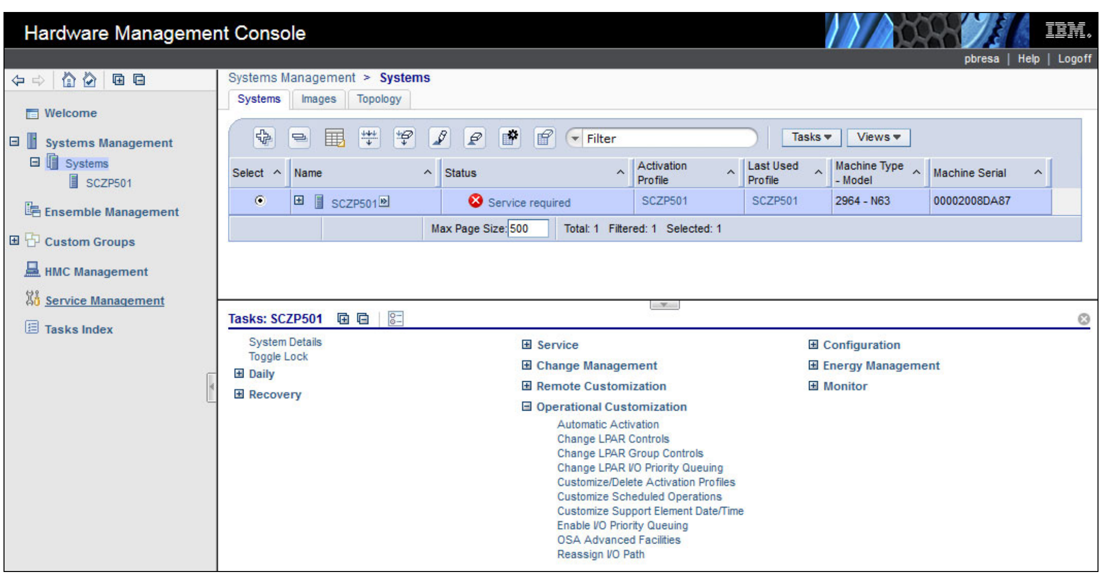

Figure 8-130 HMC - Customize Activation Profiles

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 this new TESTRESET Profile and click **Customize profile**.
8. Click the IOCDS that you updated in 8.8, “HCD/HMC: Loading the 2964 processor IOCDS” on page 495. The ACTB0PDL message is displayed (Figure 8-131).

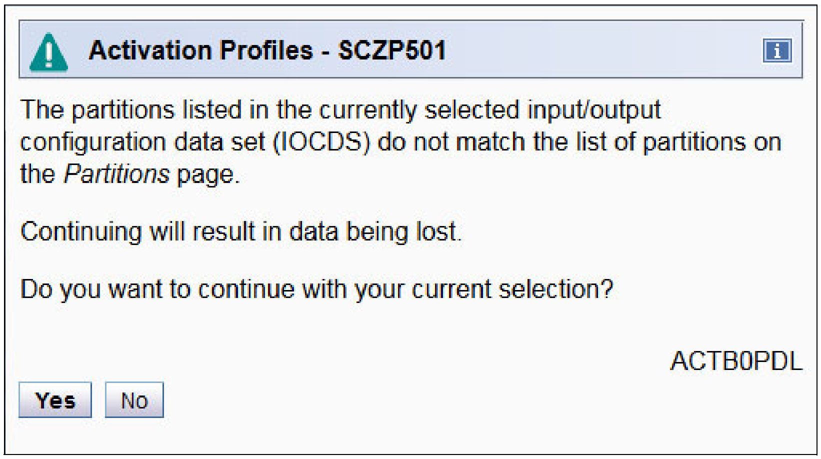

Figure 8-131 HMC - Activation Profiles: ACTB0PDL message

9. Depending on the circumstances, you can answer **Yes** or **No**. You might want to review the Partition Activation List now. For this example, we click **Yes.**
10. The HMC retrieves any Image profiles that match the LPAR names in the IOCDS that was selected. It also allows you to create new Image profiles for ones that it cannot retrieve. In the example, the last option is selected. Click **OK** (Figure 8-132).

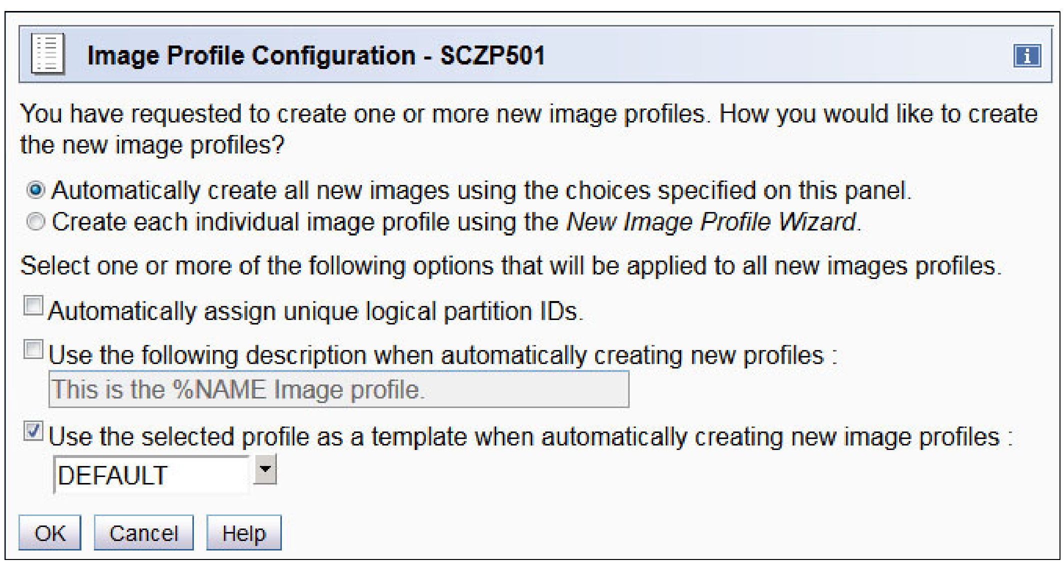

Figure 8-132 HMC - Image Profile automatic build options

11. Note the list of LPARs that were retrieved and built based on the LPARs that were defined in the selected IOCDS. Click **Save** (Figure 8-133).

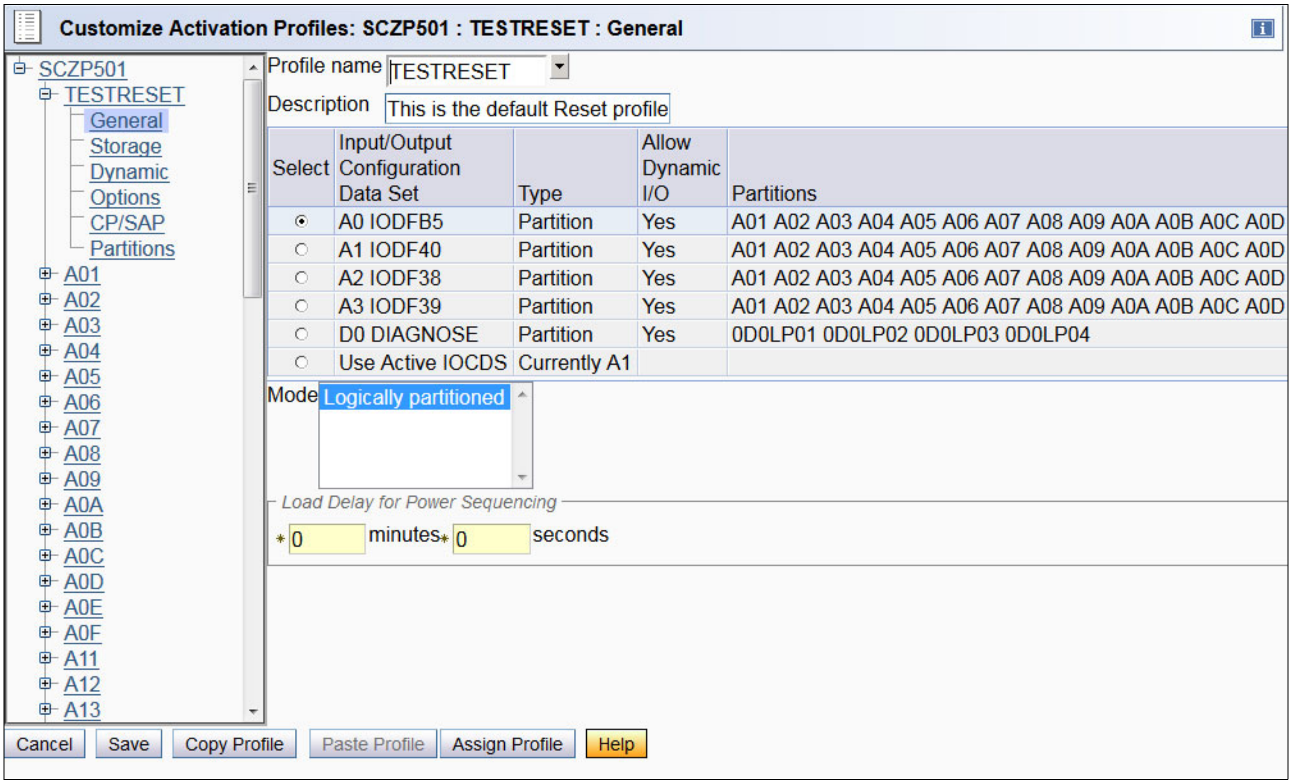

Figure 8-133 HMC - Reset and Image profile list

12. Review the items mentioned in these sections:

– 8.9.2, "Building and verifying Image Profiles" on page 508
– 8.9.5, "Server Time Protocol (STP) configuration" on page 509

8.9.2 Building and verifying Image Profiles

While still in the Reset Profile, you can review and update the Image Profile attributes.

8.9.3 Building and verifying Load Profiles

Create Load (IPL) Profiles by using the DEFAULTLOAD Load profile as a template for all the logical partitions for which you are performing an IPL on this processor.

8.9.4 Building and verifying Loadxx Members in SYS#.IPLPARM

You might need more Loadxx members defined in `SYS#.IPLPARM` for this processor.

If you used the HWNAME parameter to point to the Processor ID, update this parameter to point to the new Processor ID (in this example, from SCZP401 to SCZP501). Sometimes the LPARNAME parameter is also used in the Loadxx members and might need to be reviewed or updated.

8.9.5 Server Time Protocol (STP) configuration

Review the Server Time Protocol (STP) configuration to ensure that the correct External Time Source (ETS) was configured, and that you are connected to the correct Coordinated Time Network (CTN).

For more information about setting up your STP environment, see Chapter 9, “Server Time Protocol (STP)” on page 511.

8.9.6 Running a power-on reset (POR) of the 2964

When the 2817 or 2827 processor is upgraded with a 2964, the IBM System Service Representative will perform a power-on reset 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 power-on reset using the Reset Profile that was created in the previous step.

The 2964 is now ready to be Activated (power-on reset) using the Production Reset Profile from 8.9.1, “Building the Reset Profile for Activation and pointing to required IOCDS” on page 506.

9

Server Time Protocol (STP)

Server Time Protocol (STP) introduces the concept of Coordinated Timing Network (CTN). A CTN is a collection of servers and coupling facilities that are time-synchronized to a time value called Coordinated Server Time (CST).

This chapter describes the STP setup when a new or additional z13 installed.

This chapter includes the following topics:

- Server Time Protocol configuration overview
- Configuring the HMC as an NTP server
- Configuring a new STP-only CTN
- Adding the z13 to an existing STP-only CTN

9.1 Server Time Protocol configuration overview

STP is a CPC-wide facility that is implemented in the Licensed Internal Code (LIC) as a chargeable feature. See Figure 9-1. STP is designed to synchronize and maintain the time for various z Systems platforms coupled together in a Parallel Sysplex. z13 supports STP-only CTN.

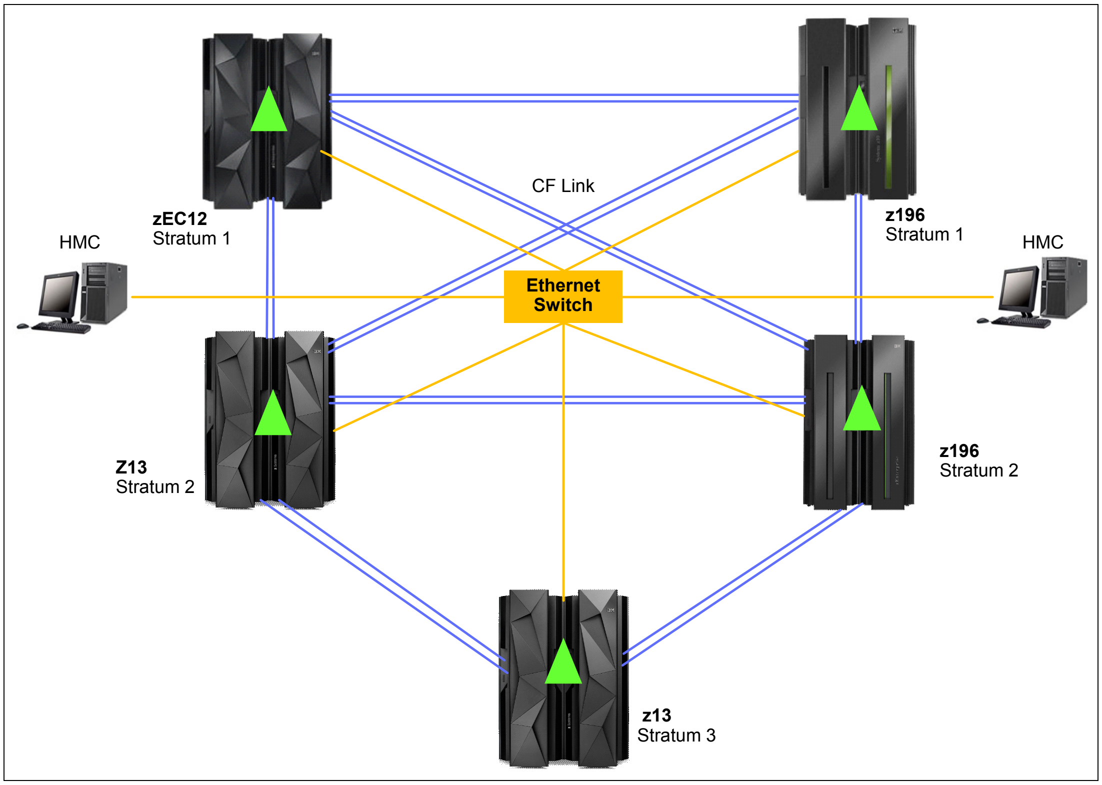

Figure 9-1   STP only CTN with z13 connectivity

> **Note:** A z13 cannot be in the same sysplex with any z Systems platform earlier than a z196. Similarly, a z13 cannot be in the same CTN with any z Systems platform earlier than a z196.

9.1.1 Network Time Protocol (NTP) server as an external time source (ETS)

To retrieve time information used to steer the CST, STP uses the Network Time Protocol (NTP) client, which runs on the Support Element (SE). An NTP client requires network connection to an NTP server.

To provide NTP data to the NTP client on the SE, you have these options:

- The HMC connected to the SE network acts as an NTP server (NTP Stratum 2 or later).
- You receive time information from a stand-alone NTP server (NTP Stratum 1 or later) through the SE network.

> **Note:** On z13, a phone modem is not supported on the HMC. Therefore, you cannot set up the HMC as ETS by using a dial-out configuration. You must set up an NTP server.

The NTP server or the NTP server with pulse per second configured as the ETS must be attached directly to the SE network. This requirement can be seen in some environments as a potential security concern. The SE network is considered in many configurations to be a private dedicated network, and must be kept as isolated as possible. Providing an HMC with the capability to act as an NTP server mitigates this security concern. The NTP server that is configured on the HMC can access another NTP server through a separate network connection. A user can define multiple NTP servers on the HMC.

The Simple Network Time Protocol (SNTP) client running on the SE can be connected to one HMC network adapter (eth0), as shown in Figure 9-2. The other HMC network adapter (eth1) is connected to the corporate network. Configuring an NTP server on the HMC can also be considered as a backup solution to provide NTP server redundancy. However, the NTP server configured on the HMC is not capable of providing a pulse per second output.

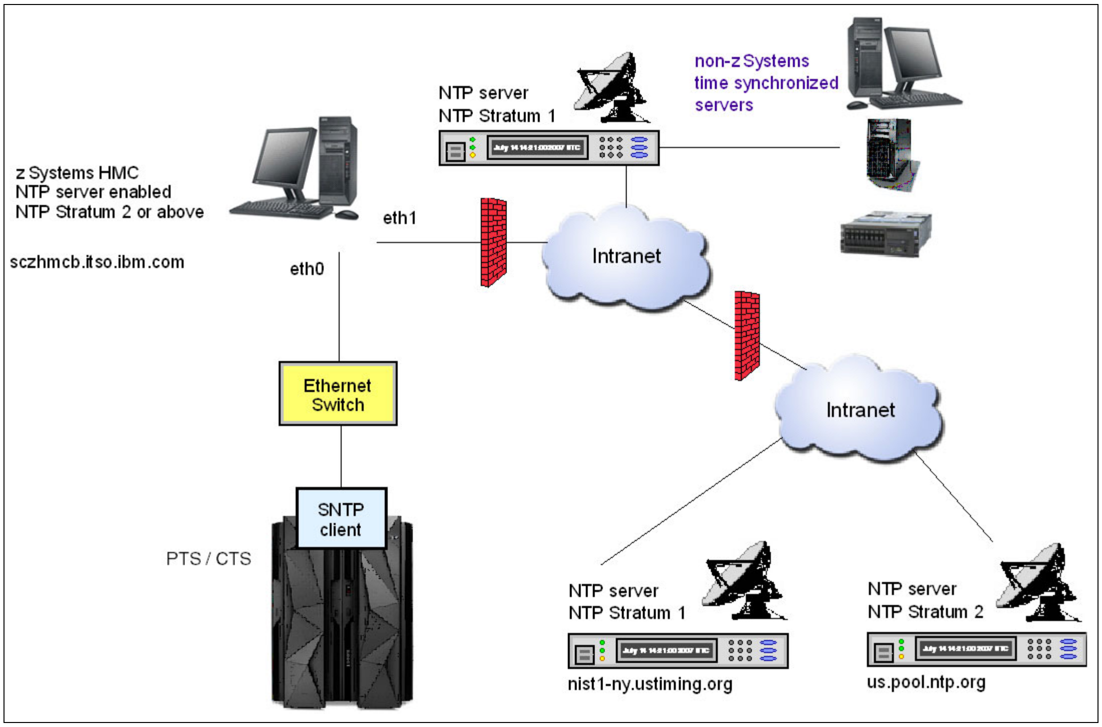

Figure 9-2 Example configuration: NTP server on the HMC

9.2 Configuring the HMC as an NTP server

This section describes how to set up the HMC as an NTP server for STP. The HMC can maintain time accuracy through an NTP server that is connected to the corporate network or available from the NTP pool on the Internet. It also addresses enabling the HMC to act as an NTP server that can be used as an ETS for the CTN.

The NTP server capability on the HMC addresses the potential security concerns that users might have about attaching NTP servers directly to the HMC/SE network. You can set up the HMC as an NTP server if you have any concerns with attaching NTP servers to your SE network directly. However, when you use the HMC as the NTP server, there is no pulse-per-second capability available.

On z13, NTP Broadband Authentication symmetric key (NTP V3-V4) and autokey (NTP V4) authentication are supported on the HMC.

For more information about NTP Broadband Authentication, see 9.2.3, “NTP Broadband Authentication for NTP server setup” on page 518.

9.2.1 Configure your HMC time source

HMC version 2.13.0 provides the following time source options:

- Network Time Protocol (NTP)
- Selected CPCs
- None

Note that if you select Network Time Protocol (NTP), the HMC will become the NTP time server for STP.

9.2.2 Procedure to set up 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 appropriate HMC. This task is available in the HMC Management section (Figure 9-3).

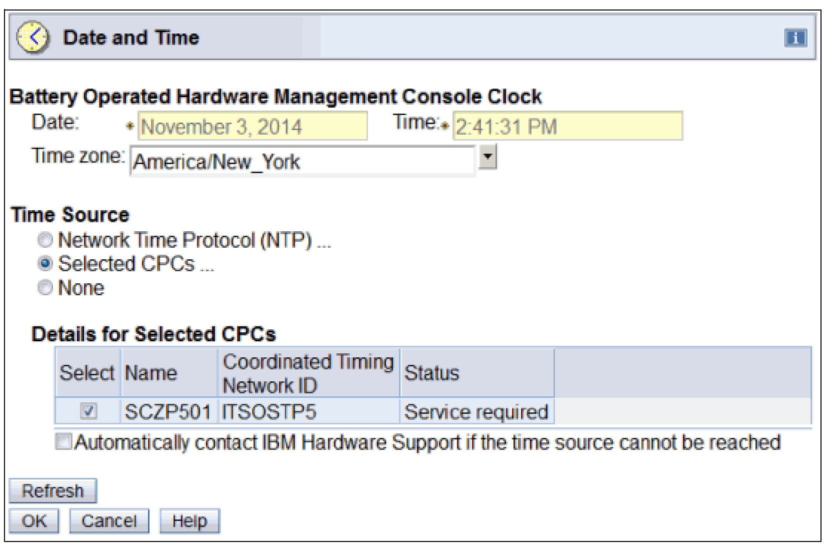

Figure 9-3 HMC: Customize Console Date and Time

2. Select **Network Time Protocol (NTP)** as the time source (Figure 9-4).

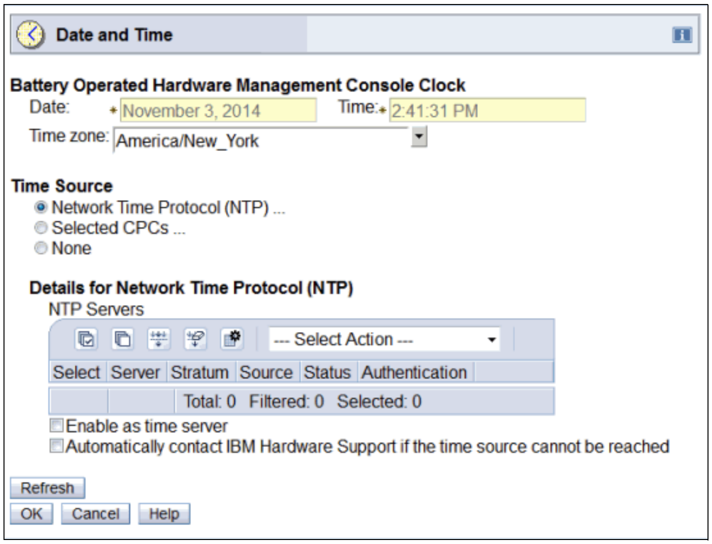

Figure 9-4 HMC: Customize Console Date and Time NTP selection

3. Select **Add Server** from the Select Action list (Figure 9-5).

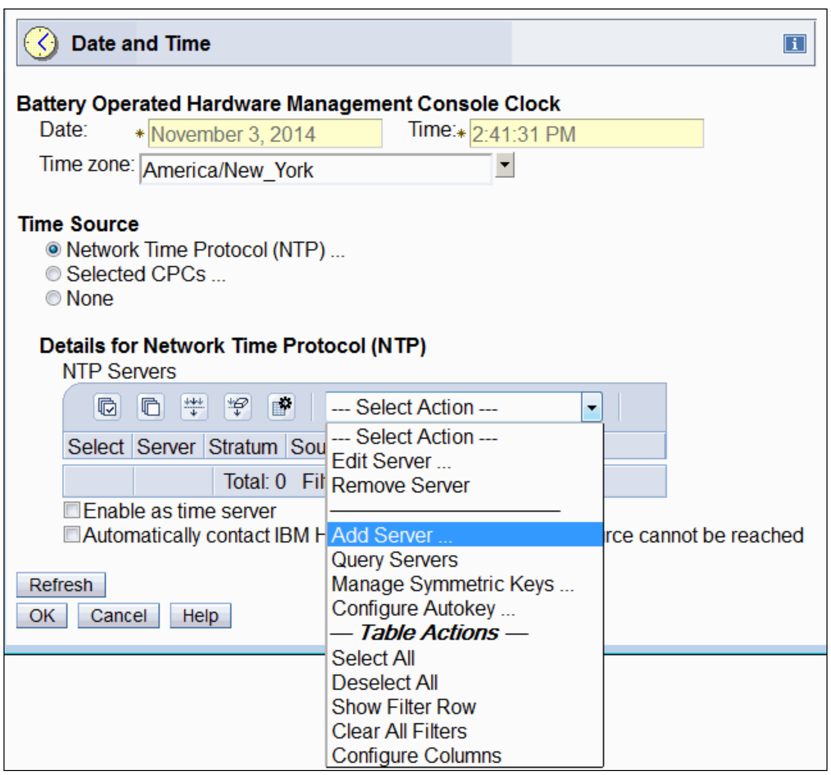

Figure 9-5 HMC: Customize Console Date and Time Add Server

4. The Add Network Time Server window opens. In this window, you can set the values in the following fields, and click **OK** when you finish (Figure 9-6 on page 516).

   – **Enter the time server host name or IP address:** This mandatory field is the IP address or host name of NTP server. For the example configuration, `0.pool.ntp.org` is the NTP server.

– **Authentication Selection:** You can define this field if you want to use NTP Broadband Authentication Support function. If you do not use NTP Broadband Authentication Support function, select **None**.

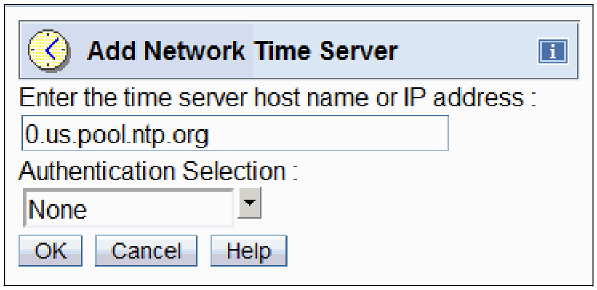

Figure 9-6 HMC: Customize Console Date and Time Add Network Time Server

Tip: Using 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.

5. During this process, the HMC tries to communicate with the NTP server you specified by using the NTP protocol to confirm whether it is reachable. If this communication successfully completes, the window shown in Figure 9-7 is displayed. You can add multiple NTP servers by repeating step 4 on page 515.

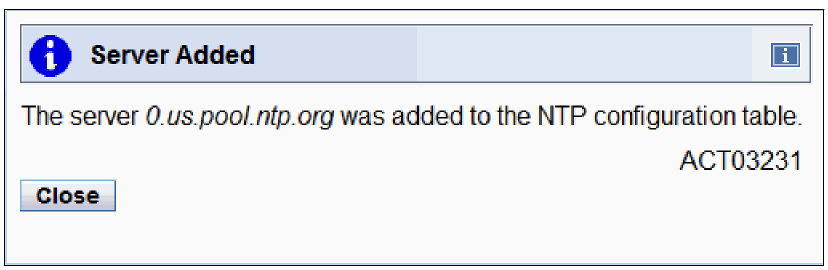

Figure 9-7 HMC: Customize Console Date and Time adding NTP server to the HMC

If the NTP server you specify is not reachable, the window shown in Figure 9-8 opens. Click **Yes** to add the NTP server for which communication failed.

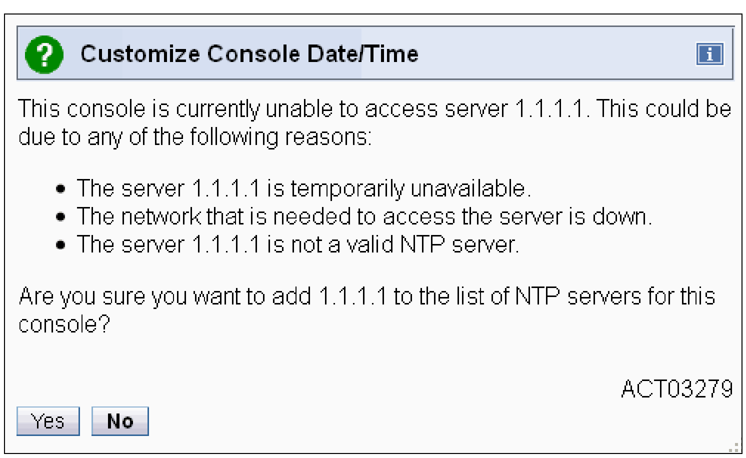

Figure 9-8 HMC: Customize Console Date and Time unreachable NTP server

6. After you add NTP Servers to the HMC, the Customize Console Date and Time window is similar to Figure 9-9. After you complete the definition of all the NTP servers, select the **Enable as time server** check box and click **OK**.

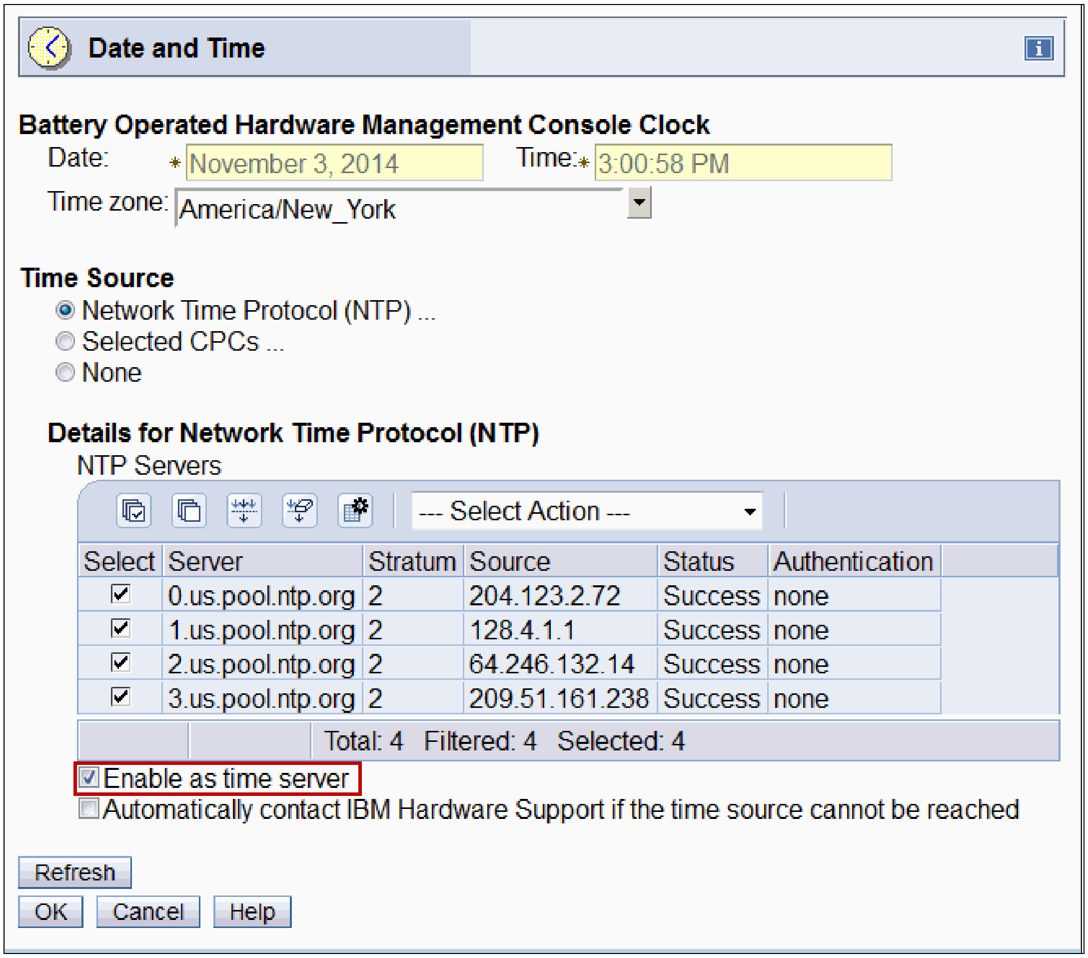

Figure 9-9 HMC: Customize Console Date and Time Configuring the NTP setting panel after you add the NTP server

7. The next panel (Figure 9-10) requests confirmation that you want to set the HMC as time server. Click **OK**.

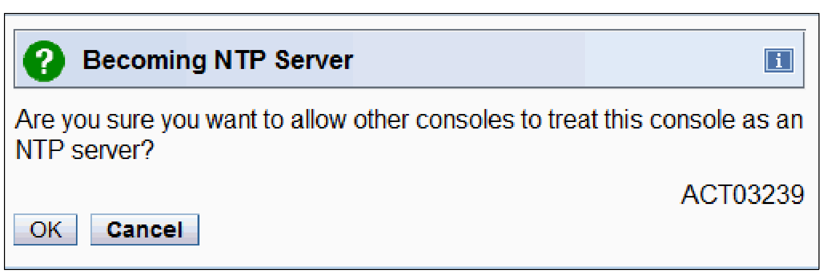

Figure 9-10 HMC: Customize Console Date and Time Becoming NTP server panel

8. After processing is completed, the next window opens (Figure 9-11 on page 518). Click **Close**. The Customize Date/Time panel closes. This is the last step of the HMC as NTP server.

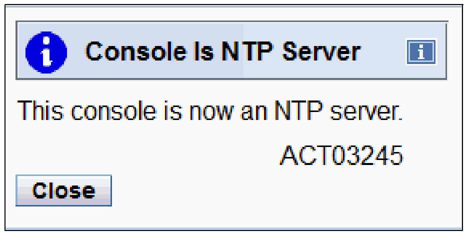

Figure 9-11 HMC: Customize Console Date and Time - Console is NTP Server panel

Consideration: If more than one NTP server is defined, you cannot specify which server is the primary server. The NTP service on the HMC will take any defined NTP server and try to contact it. If it succeeds, that server is used as the time source until either the server in question is no longer available or the console is rebooted. If it cannot communicate with that server, it tries another in the list.

The check box on the Configure NTP Setting panel is used only for actions in the **Select Action** list, and not for setting a primary or preferred NTP server.

9.2.3 NTP Broadband Authentication for NTP server setup

HMC supports Broadband Authentication. Therefore, the HMC can act as an NTP server with an increased level of security in the following situations:

- Using a proxy to access outside network

  NTP requests are User Datagram Protocol (UDP) socket packets, so they cannot pass through the proxy. The proxy must be configured as an NTP server to get to target servers on the web.

- Using firewall

  NTP requests can pass through the firewall because it uses UDP socket packets for communication. If you use firewall to access outside NTP server, use the HMC authentication to ensure untampered time stamps.

Two NTP authentication key methods are supported:

- Symmetric key (NTP V3-V4) authentication

  Symmetric key encryption uses the same key for both 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 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** on the Autokey Configuration window. Clicking this button issues the `ntp-keygen` command, which generates the specific key and certificate for this system. Autokey authentication is not available with an NAT firewall.

Setting up a Symmetric key

To set up a Symmetric key, complete these steps:

1. Open the Customize Console Date/Time window. This is the same operation described in step 1 on page 514.
2. Select **Manage Symmetric keys** from the Select Action list in Details for Network Time Protocol (NTP), as shown in Figure 9-12.

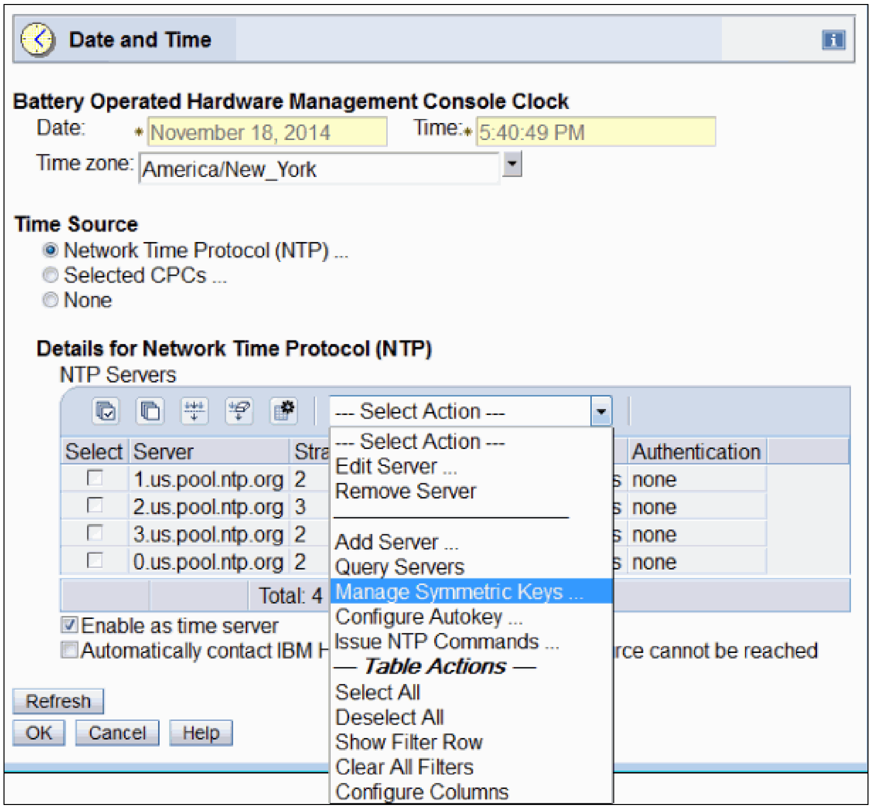

Figure 9-12 HMC: Customize Console Date and Time Manage Symmetric keys

3. The Manage Symmetric Keys window opens (Figure 9-13).
   Click **Select Action** → **Add Key**.

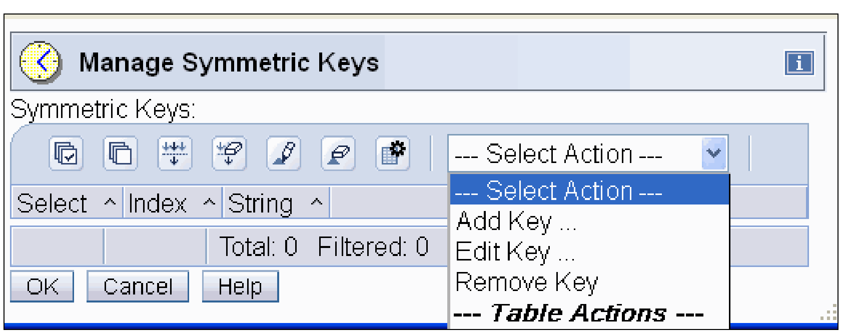

Figure 9-13 HMC: Customize Console Date and Time Add Symmetric keys

4. The Add Symmetric Key Data panel opens (Figure 9-14). These two fields require input:

   Key index A numeric value that can range from 1 to 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 less than 40 characters long, the characters can be any printable ascii character.

   Click **OK** after you enter the key data.

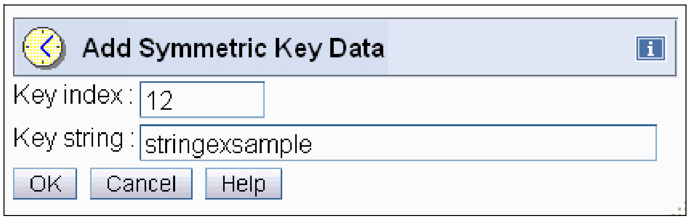

Figure 9-14 HMC: Customize Console Date and Time Symmetric key index and string

5. The Enter Symmetric Key window opens (Figure 9-15). Click **OK** to proceed.

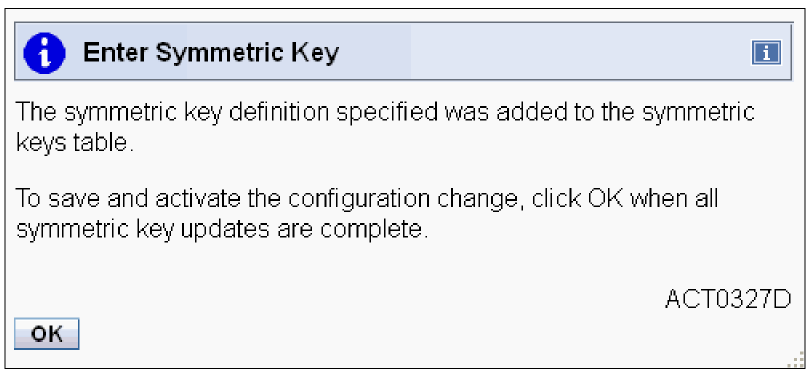

Figure 9-15 HMC: Customize Console Date and Time Save and Activate symmetric key

6. The key is displayed in the Manage Symmetric Keys window (Figure 9-16). You can enter more key data in this window by repeating the previous steps. Click **OK** to proceed if no more keys need to be added.

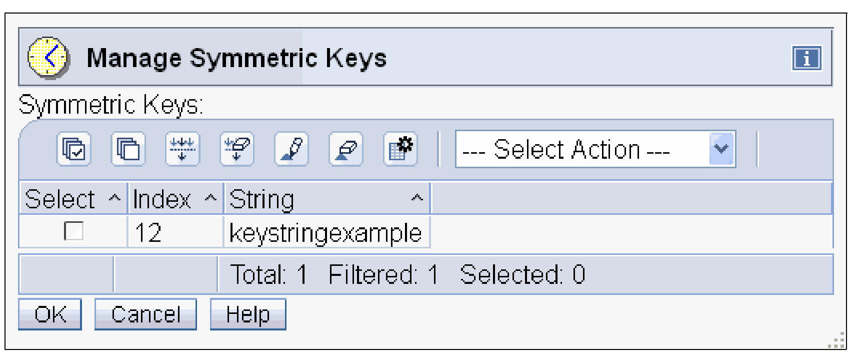

Figure 9-16 HMC: Customize Console Date and Time Saved symmetric key

7. A completion window opens (Figure 9-17). Click **OK** to return to the Configure NTP Settings window.

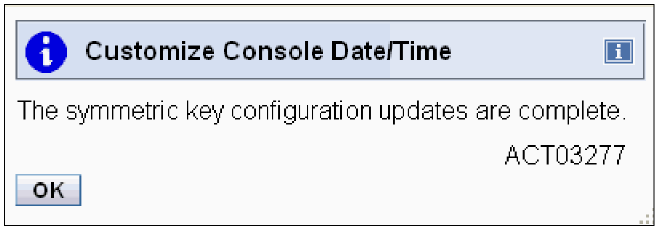

Figure 9-17 HMC: Customize Console Date and Time Completed message for symmetric key configuration

8. You can define a Symmetric Key to NTP in two ways:
   – Define a new NTP server and specify a Symmetric Key
   – Modify an existing NTP server to use a Symmetric Key

   For this example, modify NTP server 9.12.5.155 with a Symmetric Key.

9. Select the check box for the server name you want to modify and then click **Select Action** → **Edit Server** (Figure 9-18).

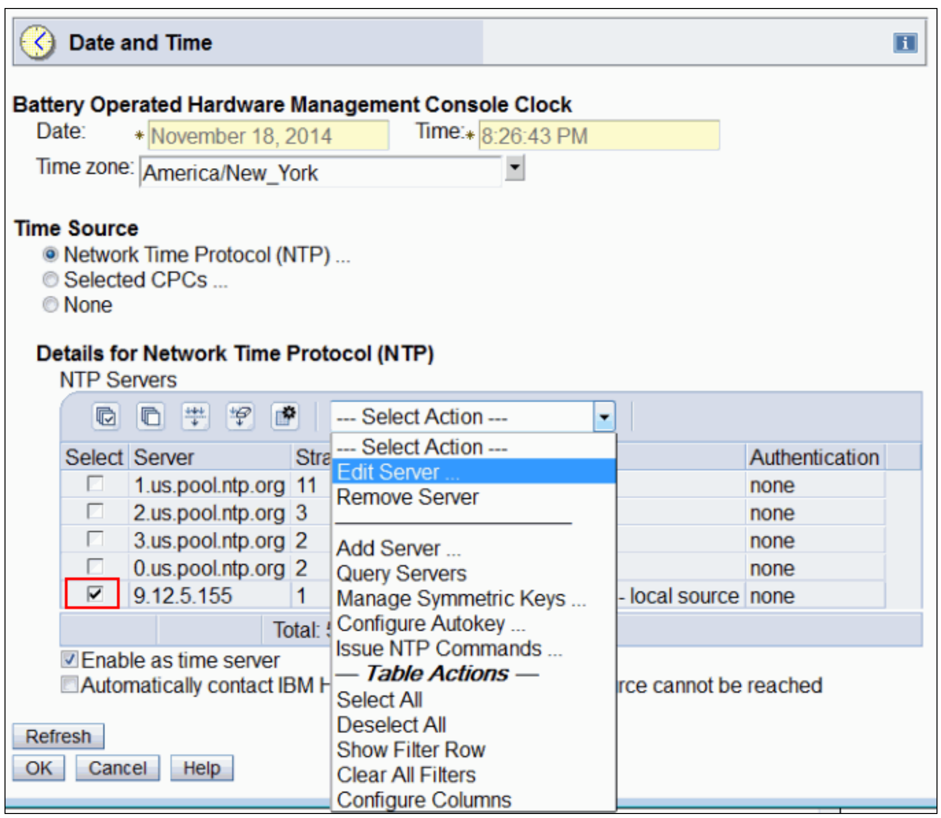

Figure 9-18 HMC: Customize Console Date and Time Server selection to edit server

10. On the Edit Network Time Server panel (Figure 9-19), select the appropriate key number in the **Symmetric Key** field. In this example, we select key **12**. Click **OK**.

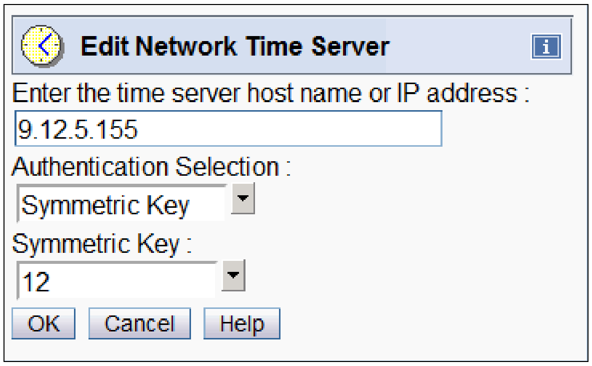

Figure 9-19 HMC: Customize Console Date and Time Adding NTP server with Authentication

11. The next window (Figure 9-20) indicates that the modification is completed. Click **Close**.

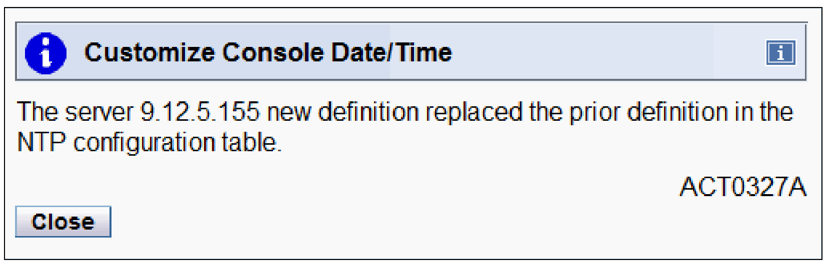

Figure 9-20 HMC: Customize Console Date and Time Completed message for adding NTP with Authentication

12. You are returned to the Date and Time window (Figure 9-21). Make sure the key that you selected is displayed in the Authentication column and the Status is `Success`.

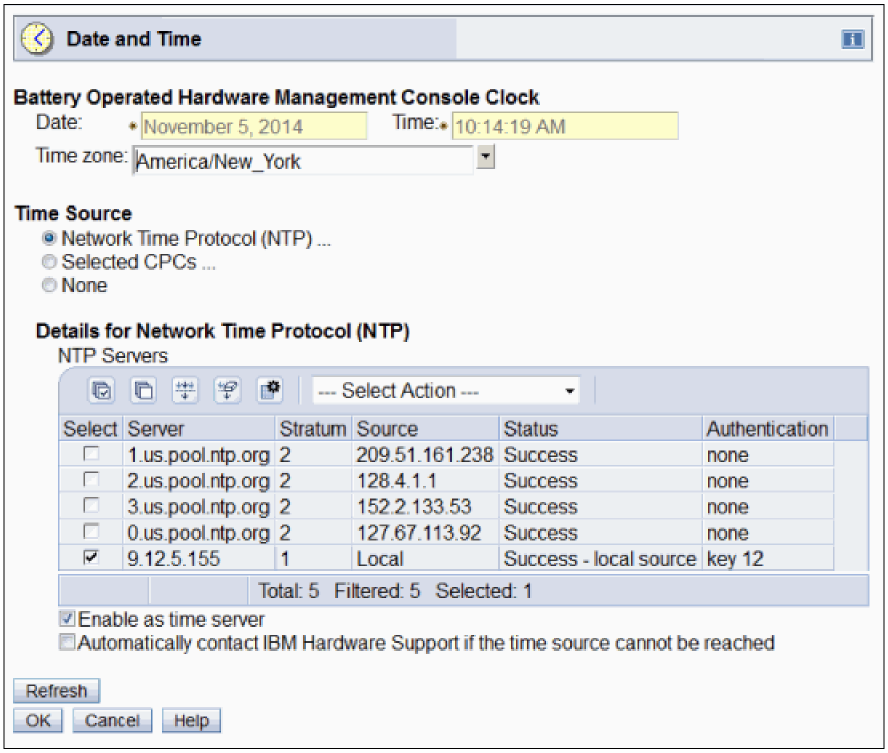

Figure 9-21 HMC: Customize Console Date and Time NTP server with Symmetric Key

Setting a Symmetric key Autokey

To set a Symmetric key Autokey, complete these steps:

1. Open the Customize Console Date/Time window. This is the same operation described in step 1 on page 514.
1. Click **Select Action** → **Configure Autokey** under the Details for Network Time Protocol (NTP) section (Figure 9-22 on page 524).

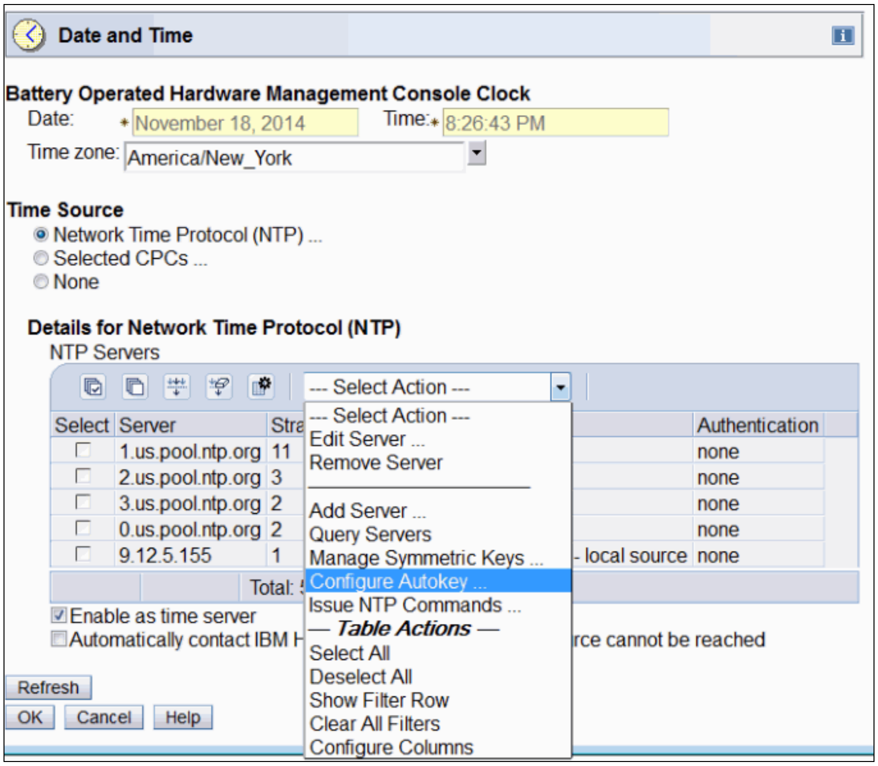

Figure 9-22 HMC: Customize Console Date and Time Configure Autokey

2. The Autokey Configuration window opens (Figure 9-23). To generate keys, click **Generate**.

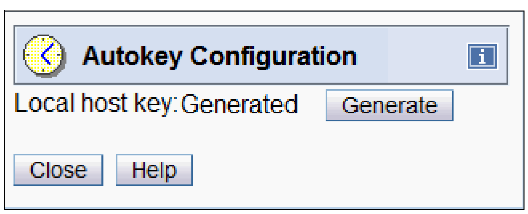

Figure 9-23 HMC: Customize Console Date and Time Generate Autokey

3. You can define an Autokey to NTP in these ways:
   - Define a new NTP server and specify an Autokey
   - Modify an existing NTP server to use an Autokey

   In this example (see Figure 9-24 on page 525), modify the existing NTP server `9.12.15.155` to use Autokey configuration.
4. Select the server name you want to modify, then click **Select Action** → **Edit Server**.
5. The Edit Network Time Server window opens (Figure 9-24 on page 525). Select **Autokey** under Authentication Selection, and then click **OK**.

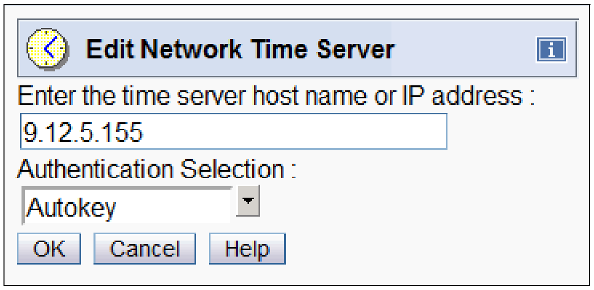

Figure 9-24 HMC: Customize Console Date and Time Edit Network Time Server for Autokey

6. The next window (Figure 9-25) indicates that the modification finished. Click **Close**.

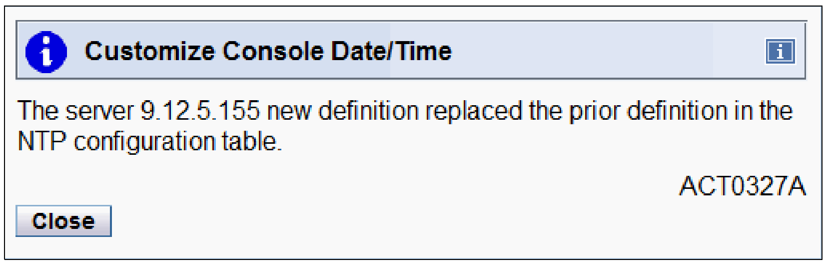

Figure 9-25 HMC: Customize Console Date and Time NTP setting change confirmation screen

7. Return to the Configure NTP Setting window. Make sure that `autokey` is listed in the Authentication column, and its status is `Success` (Figure 9-26).

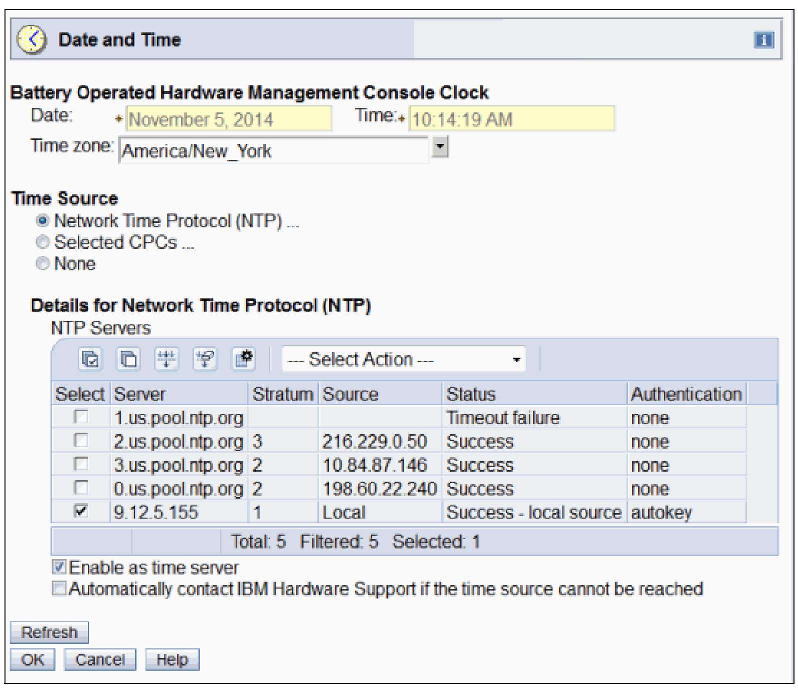

Figure 9-26 HMC: Customize Console Date and Time Successfully editing NTP server with Autokey

9.3 Configuring a new STP-only CTN

This section describes how to configure new STP-only CTN. This scenario applies to a newly installed z13 or isolated STP-only CTN for testing purpose. This example defines an STP-only CTN for newly installed z13 named SCZP501.

To configure new STP-only CTN, use these steps, which are detailed in the next sections:

1. Open System (Sysplex) Time panel
2. Set up External Time Source (ETS)
3. Define CTN ID
4. Initialize time
5. Assign preferred time server (PTS) role

Before you begin to customize STP, see these publications:

- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Implementation Guide*, SG24-7281
- *Server Time Protocol Recovery Guide*, SG24-7380

9.3.1 Open System (Sysplex) Time panel

Complete the following steps:

1. In the HMC, select **System (Sysplex) Time**, under the Configuration task (Figure 9-27).

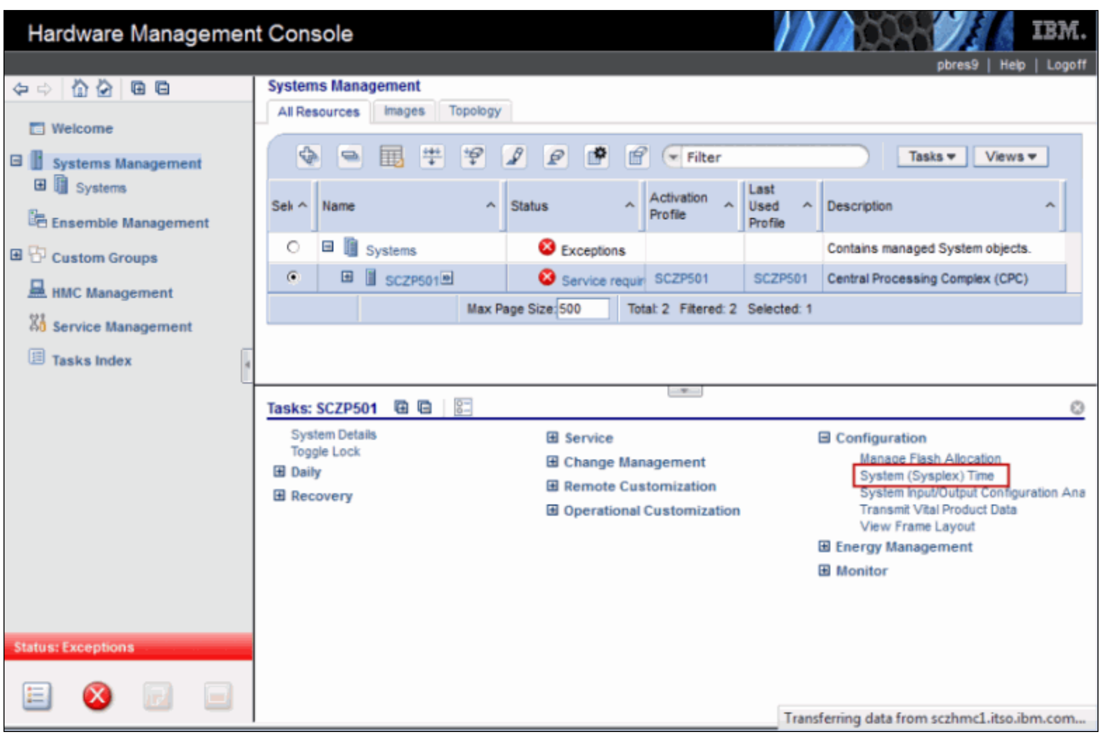

Figure 9-27 HMC: System (Sysplex) Time selection

2. Because changing System (Sysplex) time is a disruptive operation, an attention message is displayed (Figure 9-28). Read it and click **Yes** to continue.

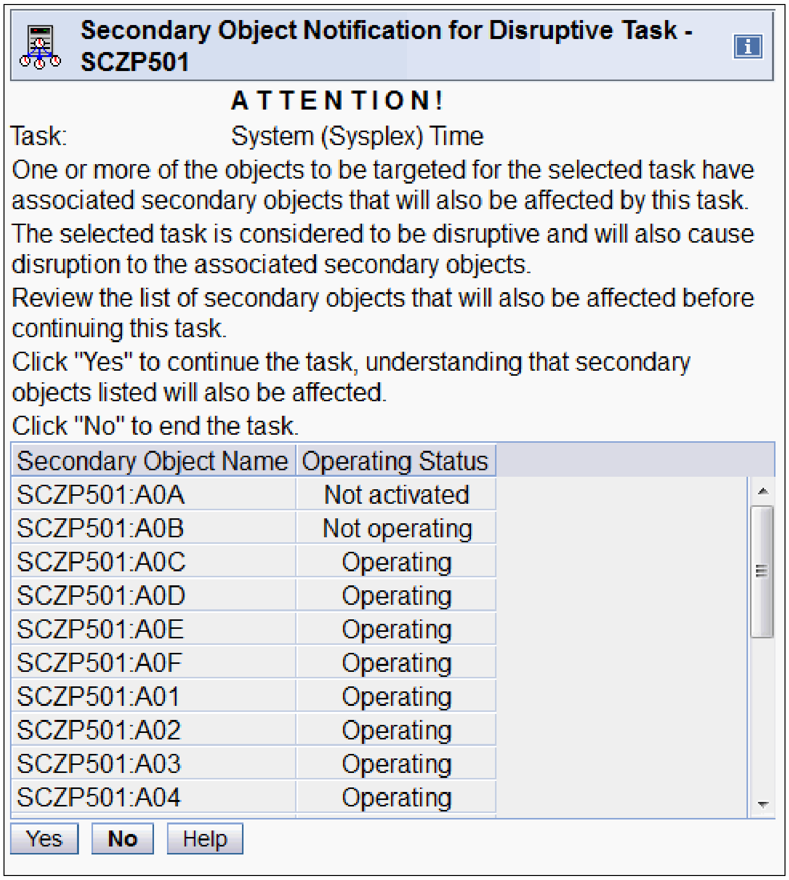

Figure 9-28 System (Sysplex) Time: Attention message

3. The System (Sysplex) Time panel opens (Figure 9-29) showing an initial status.

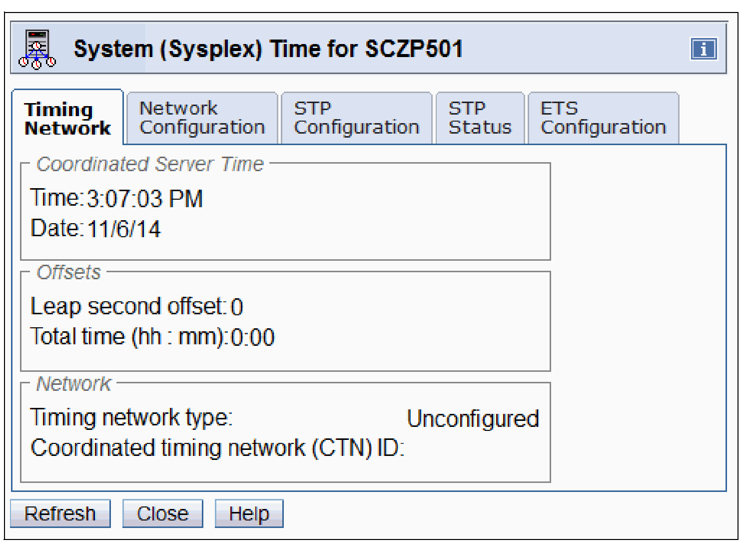

Figure 9-29 System (Sysplex) Time: Initial status

9.3.2 Set up External Time Source (ETS)

Complete the following steps:

1. Click the **ETS Configuration** tab (Figure 9-30).

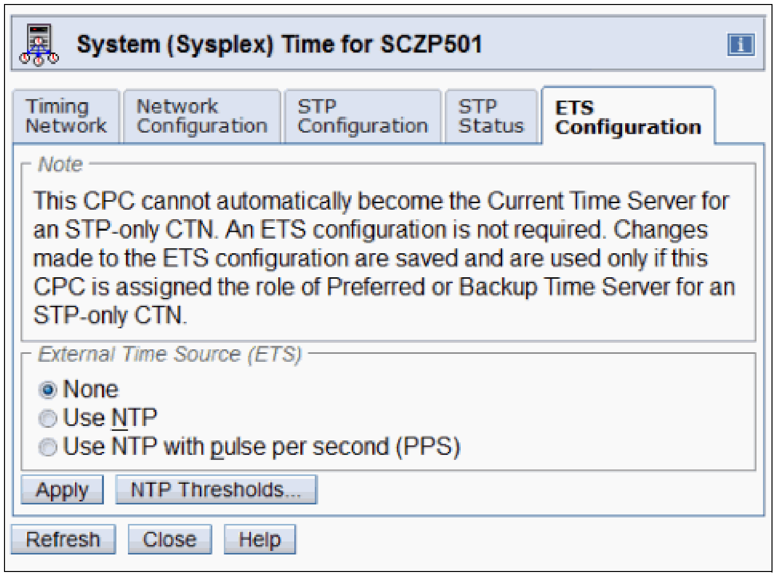

Figure 9-30 System (Sysplex) Time: ETS Configuration tab

2. Select **Use NTP** to display the NTP server information table (Figure 9-31).

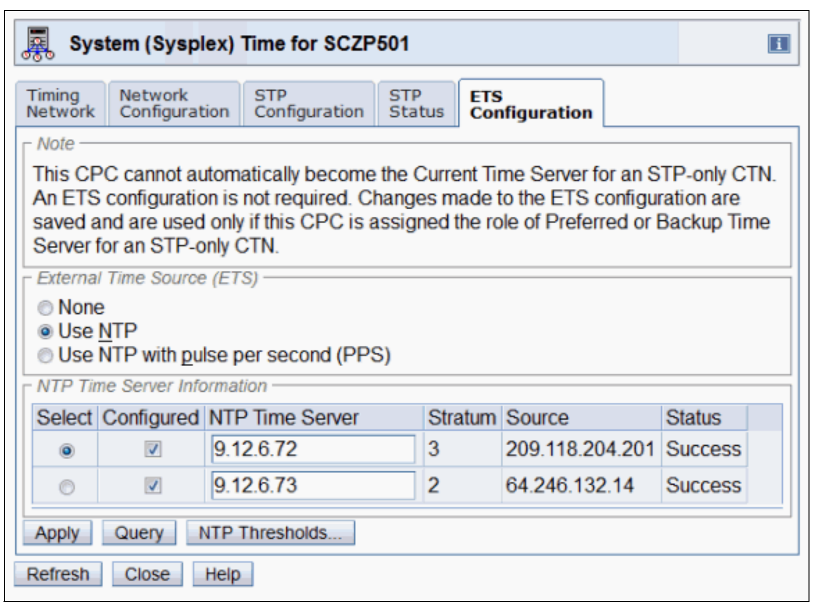

Figure 9-31 System (Sysplex) Time: two NTP server configured as ETS

3. Enter the NTP server IP address or host name in the **NTP Time Server** field. When at least one **Configured** check box is selected, the IP addresses or host names entered as the NTP server address will be used.

   If you want to use HMC as the NTP server, specify the IP address or host name of the HMC to the NTP Time Server field. If you want to use the PPS option for NTP server, select **Use NTP with pulse per second (PPS)** button in the External Time Source (ETS) section.

Up to two NTP servers can be configured. To set a preferred NTP server, select the appropriate **Select** radio button. This NTP server is called the *selected* NTP server.

4. Click **Query** to test the designated servers; the values for the Stratum level, Source, and Status table fields for the corresponding NTP server are automatically completed. The Status column displays Success if the related NTP server is accessible.
5. When you finish configuring NTP servers, click **Apply**.
6. The NTP Selection Verification panel opens (Figure 9-32). Click **Yes**.

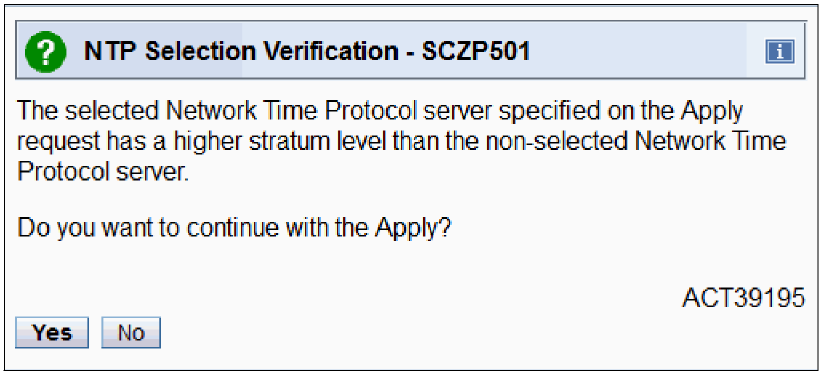

Figure 9-32 System (Sysplex) Time: NTP Selection Verification

7. The Set ETS Configuration panel (Figure 9-33) shows which selected NTP server is used for time adjustments. Click **OK** to complete the operation.

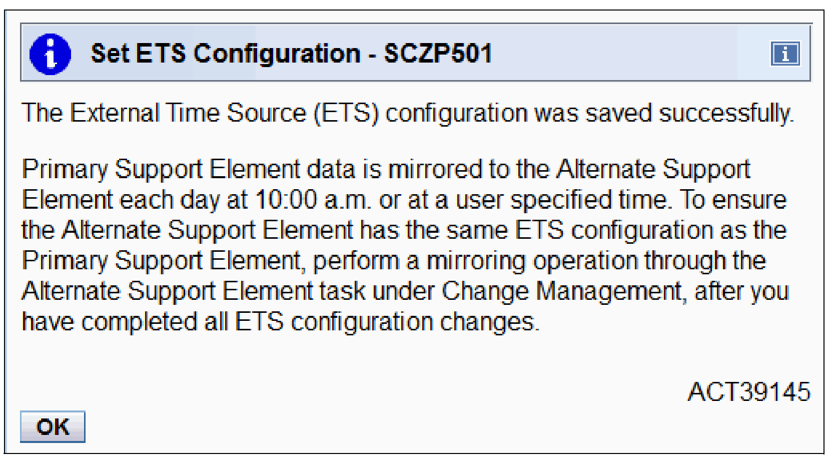

Figure 9-33 System (Sysplex) Time: Set ETS Configuration

9.3.3 Define CTN ID

The CTN ID is an identifier that is used to indicate whether the server is configured to be part of a CTN and, if so, identifies the CTN.

The CTN ID uses the format `STP Network ID - ETR Network ID`, and is the basis for the establishment of the Coordinated Timing Network.

The format of these fields is as follows:

`CTN ID = “cccccccc - xx”`

In the format, *cccccccc* is the STP Network ID and *xx* is the ETR Network ID:

- STP Network ID is case-sensitive and is 1 - 8 characters. The valid characters are as follows:

  A - Z
  a - z
  0 - 9
  - (hyphen)
  _ (underscore)

- ETR Network ID is *always empty* for an STP-only CTN.

Define the CTN ID as follows:

1. Click the **STP Configuration** tab (Figure 9-34). Enter the CTN ID. In this example, we define `ITSOSTP5` for the CTNID. Click **Apply**.

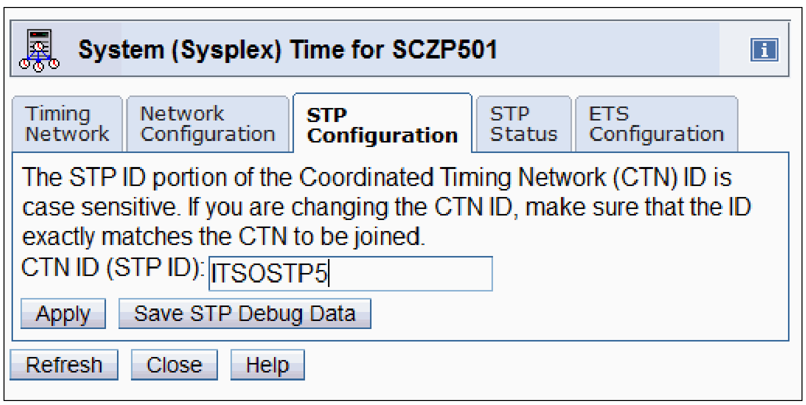

Figure 9-34 System (Sysplex) Time: STP Configuration

2. The Local Coordinated Timing Network ID Change Confirmation panel opens (Figure 9-35). Click **Yes**.

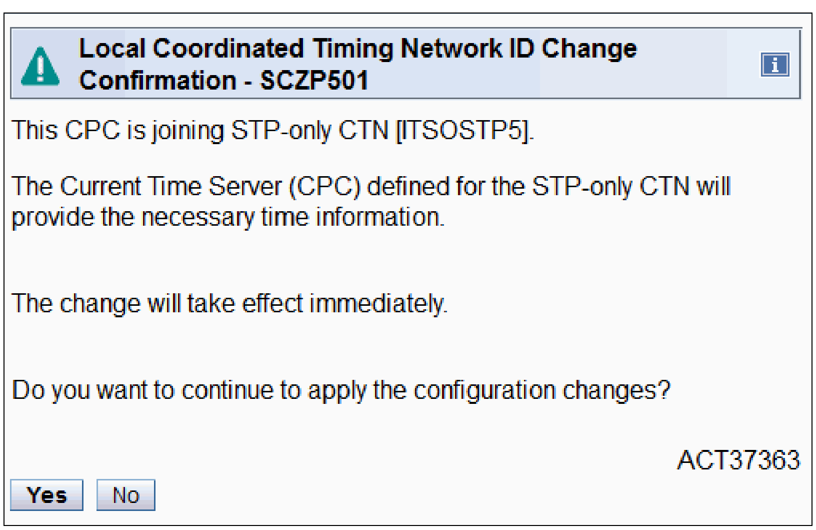

Figure 9-35 System (Sysplex) Time: Local Coordinated Timing Network ID Change Confirmation

3. When the CTNID setting is complete, the Local Coordinated Timing Network ID Change panel opens (Figure 9-36). Click **OK** to complete the operation.

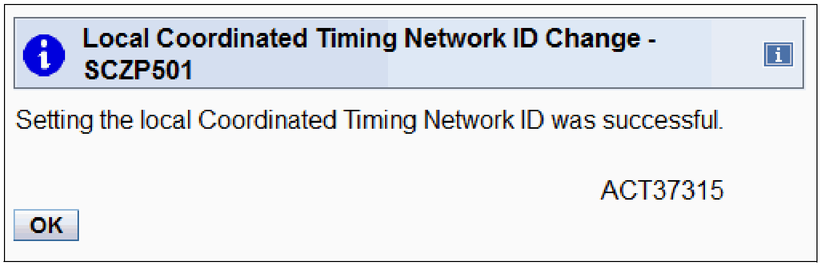

Figure 9-36 System (Sysplex) Time: Local Coordinated Timing Network ID Change

4. When CTN ID is successfully set, click the **Timing Network** tab and observe the CTN ID (Figure 9-37).

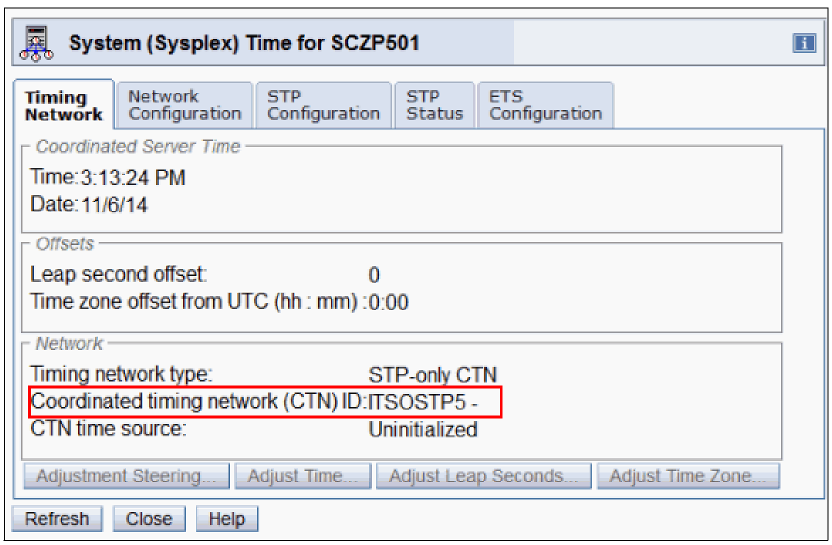

Figure 9-37 System (Sysplex) Time: Timing Network tab with CTNID

9.3.4 Initialize time

To initialize time, define three items:

- Leap second offset
- Time zone
- Date and time

Complete these steps:

1. Click the **Network Configuration** tab on the System (Sysplex) Time panel, and then click **Initialize Time** (Figure 9-38 on page 532).

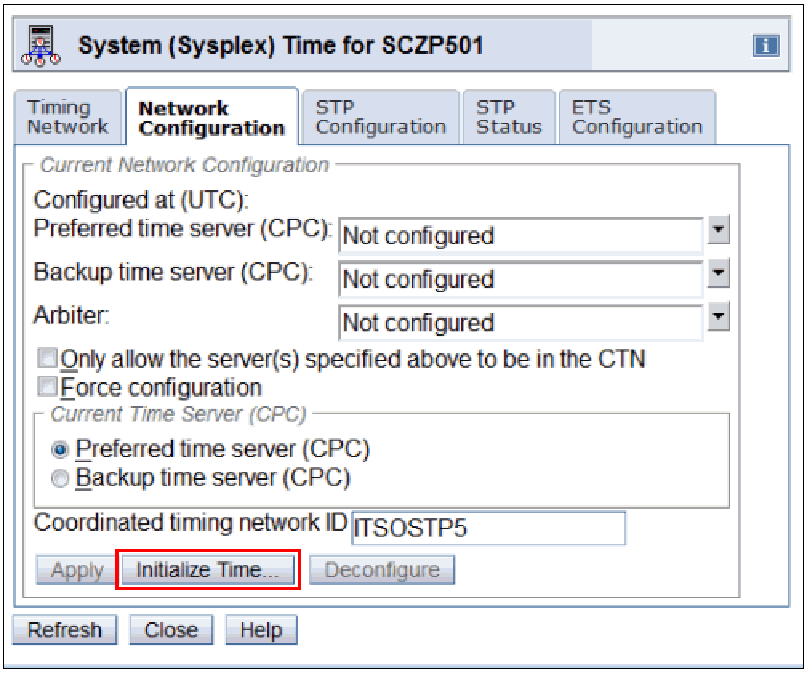

Figure 9-38 System (Sysplex) Time: Initialize Time on Network configuration

2. The Initialize Time panel opens (Figure 9-39).

   Modify the *leap second offset*:

   a. Click **Modify**, next to Leap second offset.

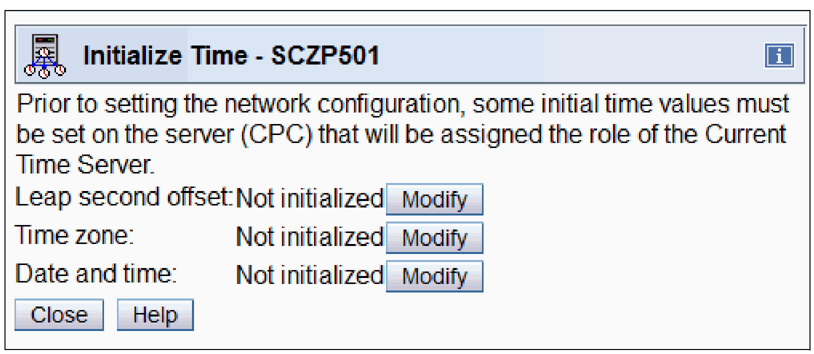

Figure 9-39 System (Sysplex) Time: Initialize Time

   b. In the Adjust Leap Second Offset panel (Figure 9-40), enter the leap second offset value. We set 25 seconds for the offset. Click **OK**.

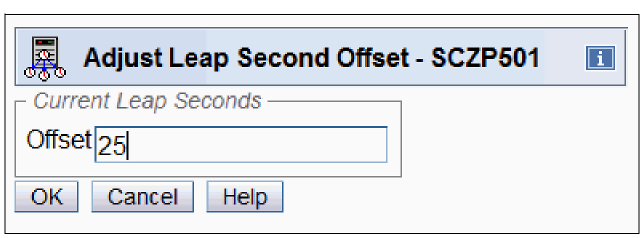

Figure 9-40 System (Sysplex) Time: Adjust Leap Second Offset panel

c. The Leap Second Offset Adjustment panel opens (Figure 9-41). Click **OK**.

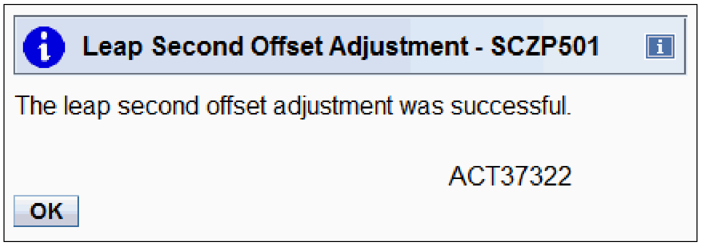

Figure 9-41 System (Sysplex) Time: Leap Second Offset Adjustment panel

Modify the *time zone*:

a. Click **Modify**, next to Time zone (Figure 9-42).

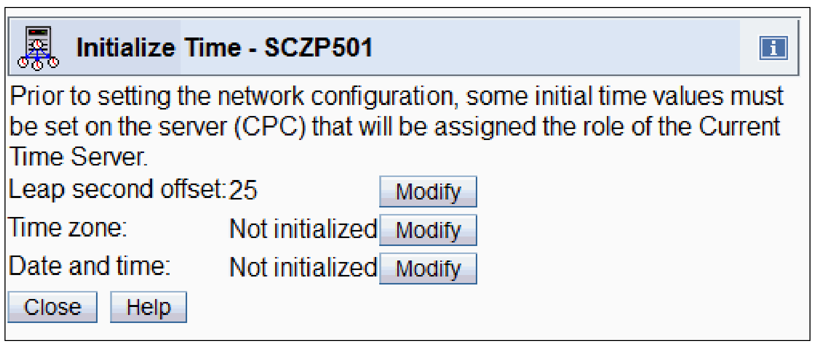

Figure 9-42 System (Sysplex) Time: Modify Time zone

b. In the Adjust Time Zone Offset panel (Figure 9-43), select a time zone (we selected **UTC-05:00 Eastern Time**) and click a daylight saving time offset option (we selected **Automatically adjust**). Click **OK**.

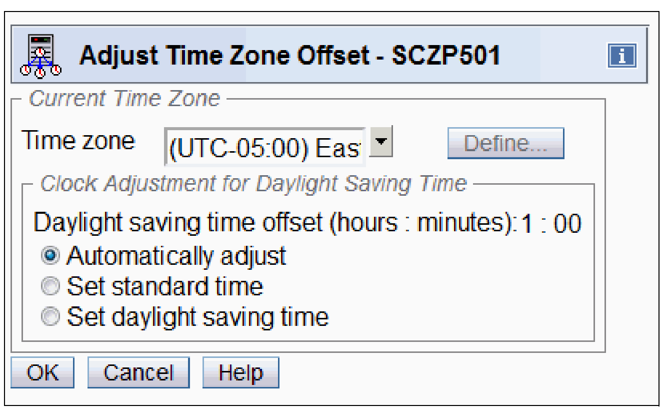

Figure 9-43 System (Sysplex) Time: Adjust Time Zone Offset Panel

c. Time Zone Algorithm panel opens (Figure 9-44 on page 534). Click **OK**.

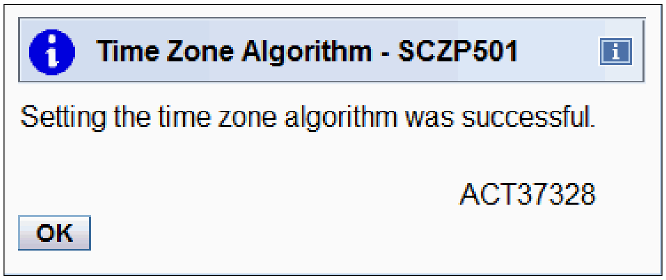

Figure 9-44 System (Sysplex) Time: Time Zone Algorithm panel

Modify the *data and time*:

a. Click **Modify**, next to Date and time (Figure 9-45).

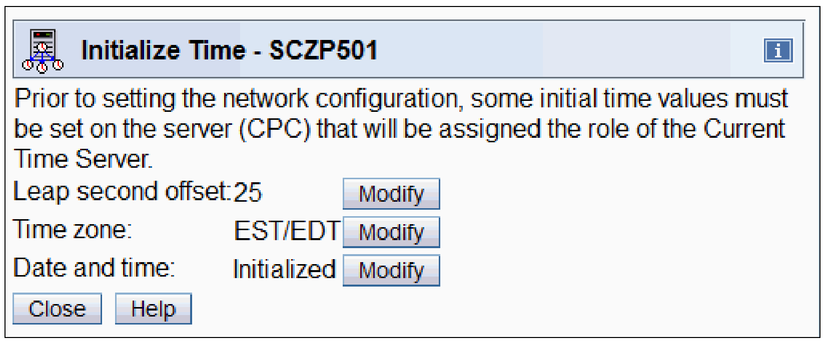

Figure 9-45 System (Sysplex) Time: Initialize Time Panel for Date and time

b. In the Set Date and Time panel (Figure 9-46), select an option. In this example, we select **Use the configured External Time Source (ETS): NTP**. Click **OK**.

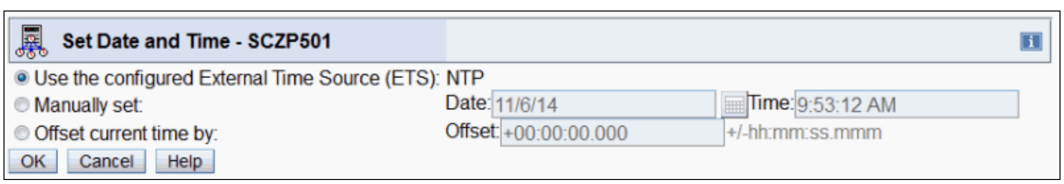

Figure 9-46 System (Sysplex) Time: Set Date and Time panel

c. A message indicates success when the operation is finished (Figure 9-47). Click **OK**.

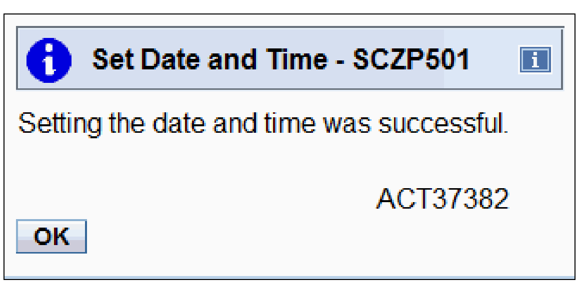

Figure 9-47 System (Sysplex) Time: Setting the date and time was successful message

3. The Initialize Time panel opens again (Figure 9-48). Verify that the settings are correct. Click **Close** to complete the process of initializing the time.

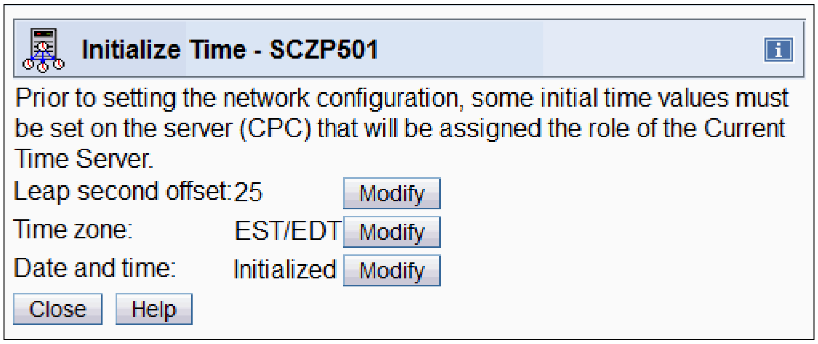

Figure 9-48 System (Sysplex) Time: Initialize Time panel after customize completed

4. The **Timing Network** tab is now similar to Figure 9-49. Verify that the Coordinated Server Time and Offsets sections show the results of your settings (Figure 9-49).

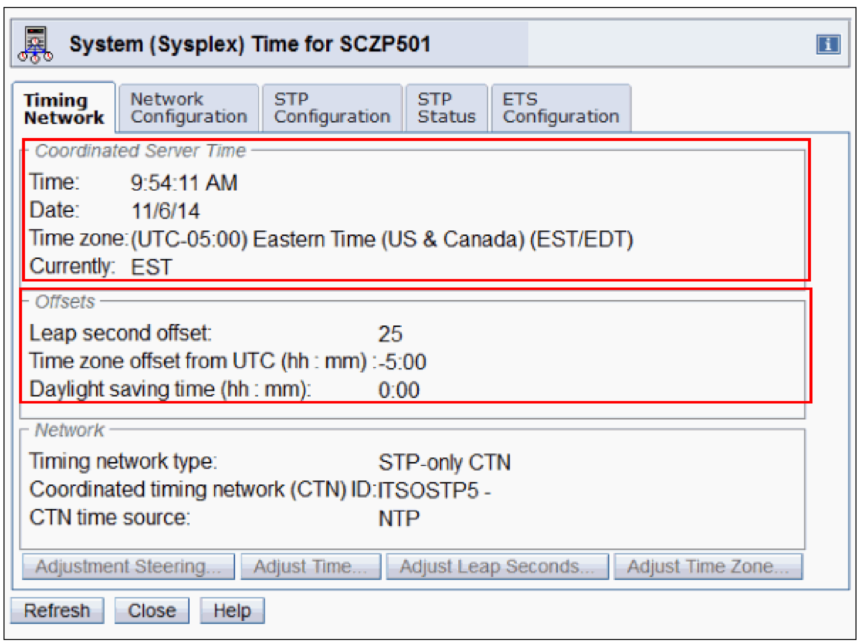

Figure 9-49 System (Sysplex) Time: Verifying the Timing Network

9.3.5 Assign preferred time server (PTS) role

The final step to set up new STP-only CTN is to set the PTS role. In this example, only one CPC will join to the STP-only CTN, so assign the new CPC as PTS.

To set the PTS role, use the following steps:

1. Click the **Network Configuration** tab on the System (Sysplex) time (Figure 9-50).

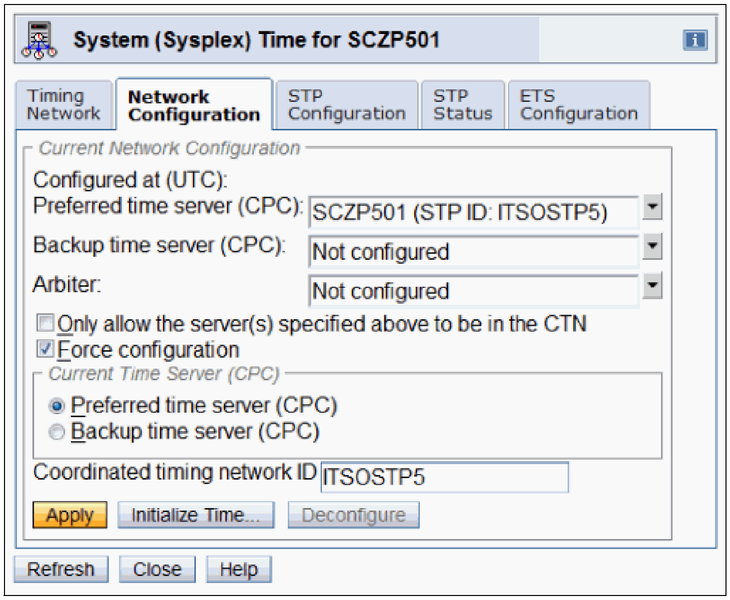

Figure 9-50 System (Sysplex) Time: Network Configuration tab: Assign PTS role

2. Select a preferred time server (CPC). In this example, we select **SCZP501** as the preferred time server (PTS).

   Also, to configure a new STP-only CTN, select the **Force configuration** check box. This option specifies whether connectivity between the PTS and other servers with a defined role are verified when a change in configuration is requested. Because this is a new STP-only CTN where no CTS is configured yet, **Force configuration** must be selected.

   > **Note:** If the **Only allow the server(s) specified above to be members of the CTN** box is selected, no other system can join this CTN. This can be removed concurrently at any time by clearing the option.

   Click **Apply** to proceed.

3. Assignment of the PTS globally changes all servers with the same CTN ID to STP timing mode. This process is confirmed by ACT37348 message. Click **Yes** (Figure 9-51 on page 537).

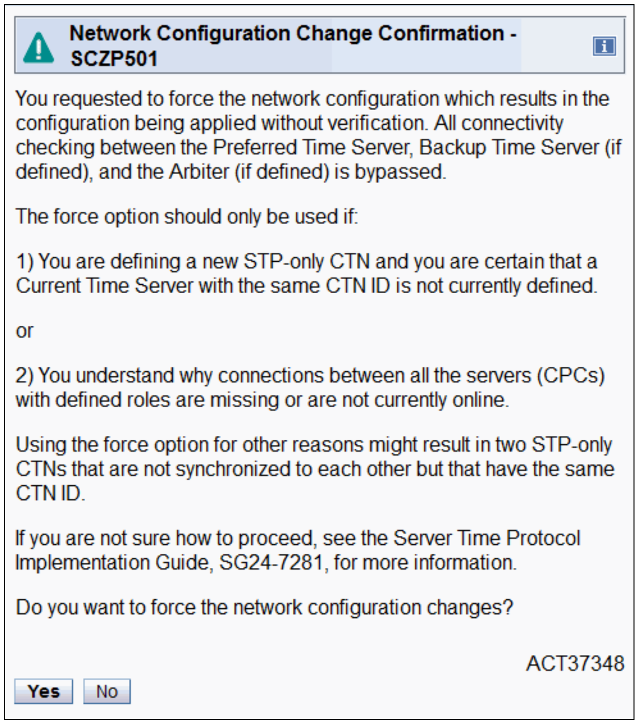

Figure 9-51 System (Sysplex) Time: Network Configuration Change Confirmation panel

4. Assigning PTS is completed. The Modify Network Configuration panel opens (Figure 9-52). Click **OK**.

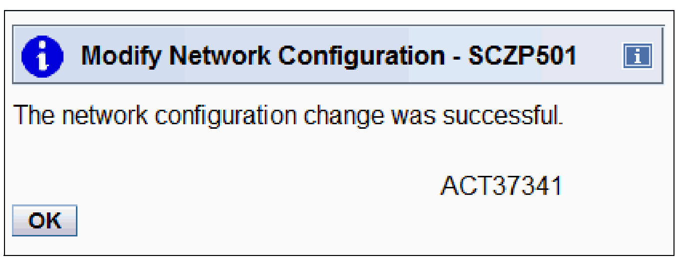

Figure 9-52 System (Sysplex) Time: The network configuration change was successful

5. All steps to configure STP-only CTN are complete. To confirm, click the **STP Status** tab and verify that following items are displayed correctly (Figure 9-53 on page 538):

Timing State:	Synchronized
Usable clock source:	Yes
Timing mode:	STP (Server Time Protocol)
Stratum level:	1

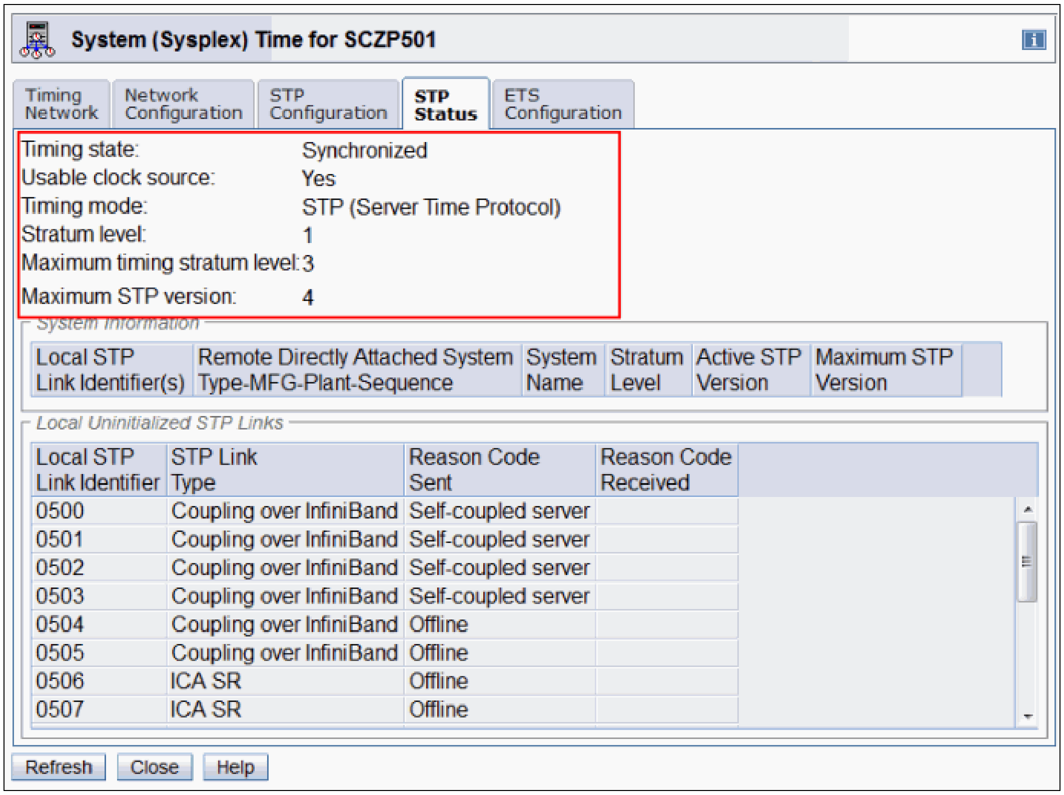

Figure 9-53 System (Sysplex) Time: STP Status panel after PTS roll is set

9.4 Adding the z13 to an existing STP-only CTN

The configuration in this example is an STP-only CTN consisting of one 2827 server. The z13 server, SCZP501, is added to this CTN.

If the STP function is enabled, the new server can be configured to join the CTN. This is always done by setting the CTN ID on the new server.

9.4.1 Setting the STP-only Network ID

Complete the following steps:

1. In the z13 HMC, open the System (Sysplex) Time panel and select the **STP Configuration** tab (Figure 9-54 on page 539).
2. Enter the CTN ID that to be joined and click **Apply**. In this example, we enter `ITSOPOK` as existing STP-only CTN. The CTN ID must match the one that is already defined to the existing STP-only CTN.

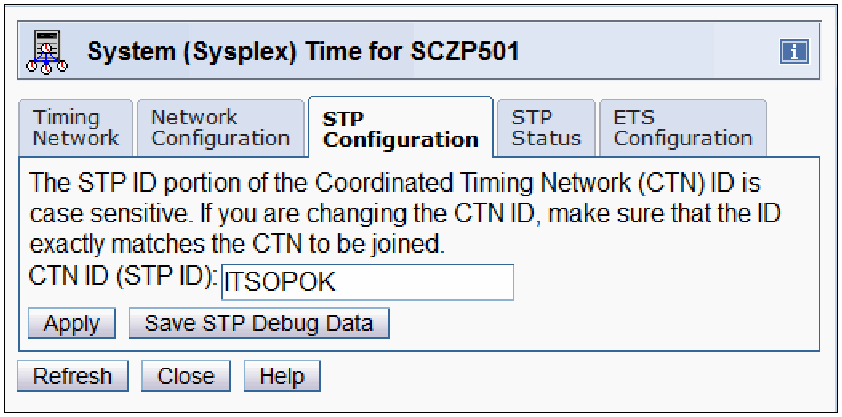

Figure 9-54 System (Sysplex) Time: STP Configuration tab defining the STP ID

A message is displayed after the CTN change is successfully completed.

9.4.2 Optional: Assign an STP role to z13

This is operation is optional. Complete the following steps If you want to assign an STP role to the z13:

1. Click the **Network Configuration** tab (Figure 9-55).
2. In the Preferred time server field, select the role that the z13 will have. You must perform this operation from PTS from SCZP401, which in our example is a PTS server on the ITSOPOK CTN ID.
3. Select the 2964 in the Backup time server field.
4. Click **Apply** to complete.

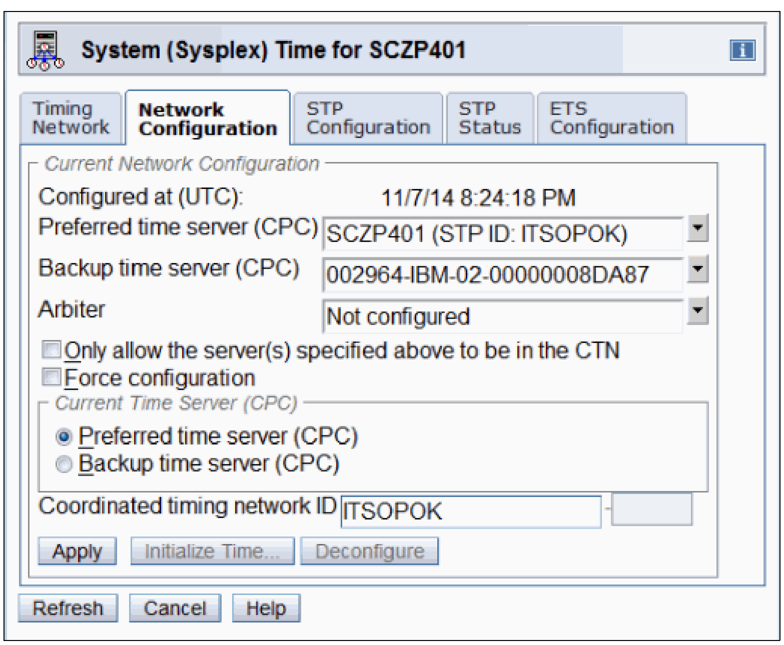

Figure 9-55 System (Sysplex) Time: Assign a new STP role

10

Crypto Express5S configuration

This chapter provides information about the configuration of the Crypto Express5S feature on an IBM z13. It covers cryptographic domains, configuration rules, and what to consider when you are planning for nondisruptive installation of cryptographic features.

The chapter includes steps for defining and configuring the Crypto Express5S feature to a logical partition.

This chapter includes the following topics:

- Crypto Express5S configuration overview
- Configuring the cryptographic feature
- Activating and deactivating cryptographic coprocessors by using ICSF

10.1 Crypto Express5S configuration overview

This section provides a brief overview of configuration rules and planning considerations.

10.1.1 Configuration rules

With zEC12 and older servers, each cryptographic coprocessor has 16 physical sets of registers or queue registers. With z13, this number is raised to 85. This corresponds to the maximum number of LPARs running on a z13, which is also 85. Each of these 85 sets belongs to a domain as follows:

- A cryptographic domain index, in the range of 0 - 84, 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 or queue registers corresponding to the domain number 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) allocated to this logical partition

The installation of the CP Assist for Cryptographic Functions (CPACF) DES/TDES enablement, Feature Code #3863, is required to use of the Crypto Express5S feature. This feature enables these 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 triple DES), and Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys.
- For key generation: Pseudo-Random Number Generator (PRNG); Random Number Generator Long (RNGL), which is 8 bytes to 8096 bytes; and Random Number Generator (RNG) with up to 4096-bit key RSA support.
- For message authentication code (MAC): Single-key MAC and double-key MAC

You can also use Secure Hash Algorithm-1 (SHA-1) 160-bit, or Secure Hash Algorithm-2 (SHA-2) for 224-, 256-, 384-, and 512-bit support. SHA-1 and SHA-2 are shipped enabled on all zEC12, and do not require an enablement feature.

The total number of Crypto Express5S features cannot exceed 16 per z13. The initial order for Crypto Express5S is two features (two PCIe adapters for Crypto Express5S). After the initial order, the minimum order is one feature.

Each Crypto Express5S feature contains one PCI-X adapter. The adapter can be in the following configurations:

- IBM Enterprise Common Cryptographic Architecture (CCA) Coprocessor (CEX5C)
- IBM Enterprise Public Key Cryptography Standards#11 (PKCS) Coprocessor (CEX5P)
- IBM Crypto Express5S Accelerator (CEX5A)

During the feature installation, the PCI-X adapter is configured by default as the CCA coprocessor.

The configuration of the Crypto Express5S adapter as EP11 coprocessor requires a Trusted Key Entry (TKE) workstation Hardware 8.0 (FC 0847) with TKE 8.0 Licensed Internal Code (FC 0877).

The Crypto Express5S feature does not use CHPIDs from the channel subsystem pool. However, the Crypto Express5S feature requires one slot in a PCIe I/O drawer, and one PCHID for each PCIe cryptographic adapter.

Table 10-1 summarizes the cryptographic feature codes for z13.

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-224, SHA-256, SHA-384, and SHA-512) and Crypto Express features.
0890	Crypto Express5S feature: ▶ A maximum of sixteen (16) features can be ordered (minimum of 2 adapters). This is an optional feature, and each feature contains one PCI Express cryptographic adapters. This feature is only supported on z13. ▶ The maximum of all features 0890 and 0865 in a z13 also is sixteen (16).
0847	Trusted Key Entry (TKE) workstation Hardware 8.0: ▶ This feature is optional. The TKE 0847 provides a basic key management (key identification, exchange, separation, update, backup), and security administration. The TKE workstation has one Ethernet port, and supports connectivity to an Ethernet local area network (LAN) operating at 10, 100, or 1000 Mbps. ▶ This TKE feature contains the new 4767 crypto adapter and is needed for TKE 8.0 LIC (FC 0877).
0842	Trusted Key Entry (TKE) workstation Hardware 7.3: ▶ This feature is optional. The TKE 0842 provides a basic key management (key identification, exchange, separation, update, backup), and security administration. The TKE workstation has one Ethernet port, and supports connectivity to an Ethernet local area network (LAN) operating at 10, 100, or 1000 Mbps. ▶ On z13, this feature can only be carried forward from a zEC12 or a zBC12. It can only hold TKE 7.3 LIC (FC 0872), not the new TKE 8.0 LIC (FC 0877). An upgrade to FC 0847 is available.
0877	TKE 8.0 Licensed Internal Code (TKE 8.0 LIC): ▶ Requires Trusted Key Entry workstation feature code 0847. ▶ Is required to support CEX5P, and also for CEX4P when using more than 16 domains. ▶ Can also be used to control Crypto Express features on zEC12, zBC12 and earlier z Systems platforms. ▶ Is shipped installed on the TKE Workstation Hardware 8.0 (FC 0847).
0872	TKE 7.3 Licensed Internal Code (TKE 7.3 LIC): ▶ Requires Trusted Key Entry workstation feature code 0842. ▶ Is required to support CEX4P (however, for CEX5P and also for CEX4P with more than 16 domains used, FC 0877 is required). The 7.3 LIC can also be used to control Crypto Express features on zEC12, zBC12 and earlier z Systems platforms. ▶ On z13, this feature can only be carried forward from a zEC12 or a zBC12, together with the old TKE workstation Hardware 7.3 (FC 0842). An upgrade to FC 0877 is available.

Feature code	Description
0885 (carry forward) 0891	TKE Smart Card Reader: ► Access to information in the smart card is protected by a personal identification number (PIN). One (1) feature code includes two Smart Card Readers, two cables to connect to the TKE workstation, and 20 smart cards. Smart card part 45D3398 or 74Y0551 may be used on TKE 8.0. A new card 00JA710 will be released because of the end of life of 74Y0551.
0884 (carry forward) 0892	TKE additional smart cards: ► When one feature code is ordered, 10 smart cards are shipped. Order increment is one up to 99 (990 blank smart cards). Smart cards 45D3398 and 74Y0551 may be used on TKE 8.0. A new card 00JA710 will be released due to the end of life of 74Y0551.

10.1.2 Configuration planning

The z13 always operates in LPAR mode. The concept of *dedicated coprocessor* does not apply to the PCI-X adapter. A PCI-X adapter, whether configured as coprocessor or accelerator, is made available to logical partitions as directed by the domain assignment and the candidate list. This occurs regardless of the shared or dedicated status that is given to the central processors in the partition.

The z13 allows for up to 85 logical partitions to be active concurrently.

Each PCI-X adapter on a Crypto Express5S feature supports 85 domains, whether it is configured as a Crypto Express5S coprocessor or a Crypto Express5S accelerator.

For availability reasons the minimum configuration consists of two Crypto Express5S, features so that every potential logical partition can have access to two cryptographic adapters.

More Crypto Express5S features might be needed to satisfy application performance and availability requirements:

- For availability, spread assignment of multiple PCI-X 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 PCI-X adapter that is configured as a Crypto Express5S coprocessor creates an application single point of failure. This point of failure exists because RSA-retained private keys are not copied or backed up.
- There is an intrusion latch within the PCI-X 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 PCI-X 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 PCI-X adapter then can be enabled from the TKE and normal operations can resume.

  Plan the definition of domain indexes and cryptographic coprocessor numbers in the Candidate list for each 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 “Change LPAR Cryptographic Controls function” on page 552.

- Crypto Express5S features can be installed concurrently when all physical requirements are fulfilled. Dynamically enabling a new PCI-X adapter to a partition requires these configurations:
  - At least one usage domain index be 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 usage domain index specified 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 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 Express5S features are assigned coprocessor numbers sequentially during the power-on-reset that follows the installation.

  However, when a Crypto Express5S 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 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 PCI-X 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 PCI-X adapter numbers

	Domain index 0	**Domain index 1**	**Domain index 2**	**.../...**	**Domain index 84**
PCI-X adapter 0	LP00 LP02	LP04	LP05		
PCI-X adapter 1	LP01 LP02				
PCI-X adapter 2					
…					
…					
PCI-X adapter 13					
PCI-X adapter 14					
PCI-X adapter 15					

Table 10-2 on page 545 illustrates these characteristics:

- Logical partitions LP00 and LP01 use domain 0 (zero), but are assigned different PCI-X adapters. There is no conflict. 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 PCI-X cryptographic adapter 0 (zero), so there is no conflict. The combination of domain number and cryptographic coprocessor number is unique across partitions.

> **Important:** Any combination of PCI-X 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 Express5S feature for z Systems, see *IBM z13 Technical Guide*, SG24-8251.

10.2 Configuring the cryptographic feature

This section provides steps for configuring Crypto Express5S for z13.

The z13 operates only in LPAR mode. For each logical partition that requires access to a PCI-X adapter, configured as either an accelerator or a coprocessor, the required information must be defined in the partition image profile. This ensures the correct use of the cryptographic features when the associated partition is activated.

Concurrent changes to the Crypto Express5S features and controls when the partition is already activated are provided by special functions at the Support Element (SE).

10.2.1 Checking whether CPACF DES/TDES enablement feature is installed

The z13 DES/TDES (FC 3863) enables the DES and TDES algorithms on the CPACF. It is a prerequisite for using the Crypto Express5S 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. Open the System details menu of the CPC at the SE workplace. The system details window opens (Figure 10-1 on page 547).

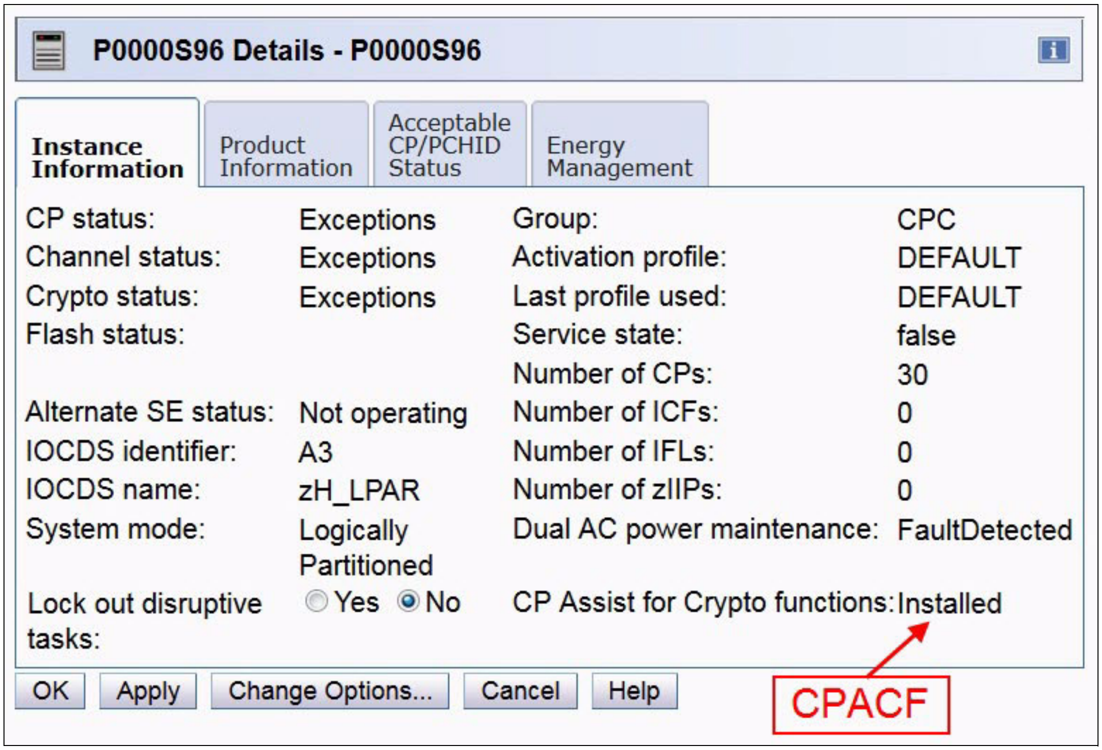

Figure 10-1 System details: CPACF installed

2. On the **Instance Information** tab, 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 feature code 3863 is enabled.
   - If the window shows the message `CP Assist for Crypto Functions: Not installed`, feature code 3863 is not installed. You can still customize the partition image profiles, but the cryptographic functions do not operate.

10.2.2 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

This task is accomplished by using the Customize/Delete Activation Profile task, which is in the Operational Customization Group, either from the HMC or from the SE. Modify the cryptographic initial definition from the Crypto option in the image profile, as shown in Figure 10-2 on page 548. 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 Support Element, 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. More information about using this function is in “Change LPAR Cryptographic Controls function” on page 552.

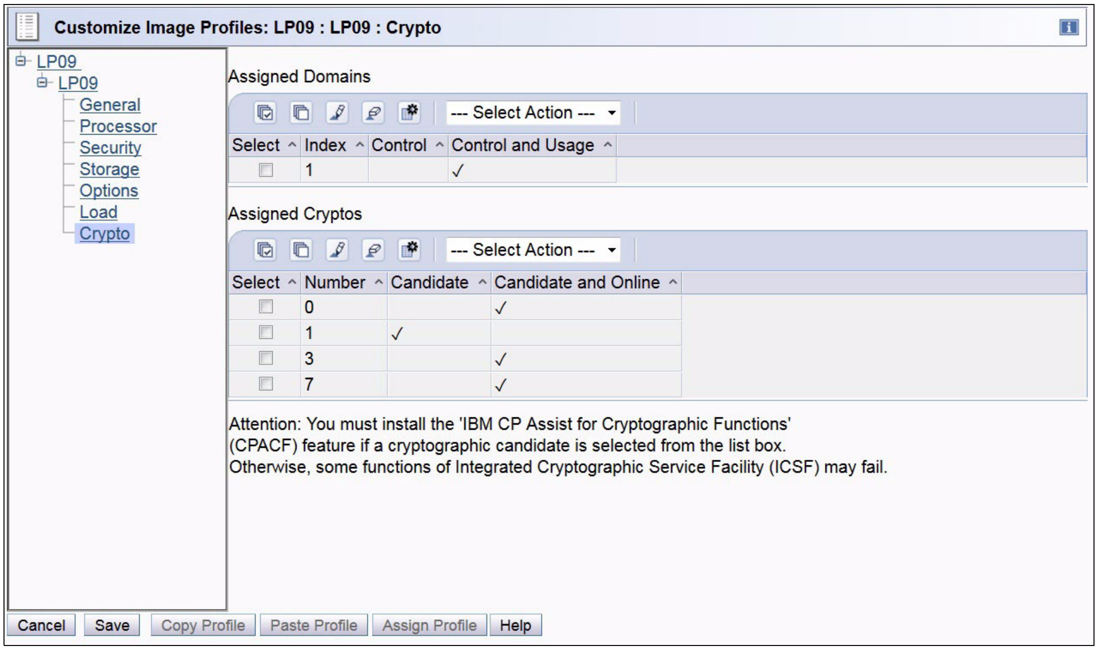

Figure 10-2 Customize Image Profiles: Crypto

The cryptographic resource definitions have 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 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 PCI-X 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 PCI-X adapters to be accessed by this partition.

No error condition is reported when a cryptographic coprocessor number, selected in the partition candidate list, is not available to the partition when the partition is activated, either because it is configured off or not installed. The cryptographic coprocessor number is ignored and the activation process continues.

If the cryptographic coprocessor number and usage domain index combination for the coprocessor that is selected is already in use by another active logical partition, the activation of the logical partition fails (Figure 10-3). 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 (Figure 10-5 on page 552). Resolve the error based on the collected data by assigning a unique combination of PCI-X adapter number and usage domain index number.

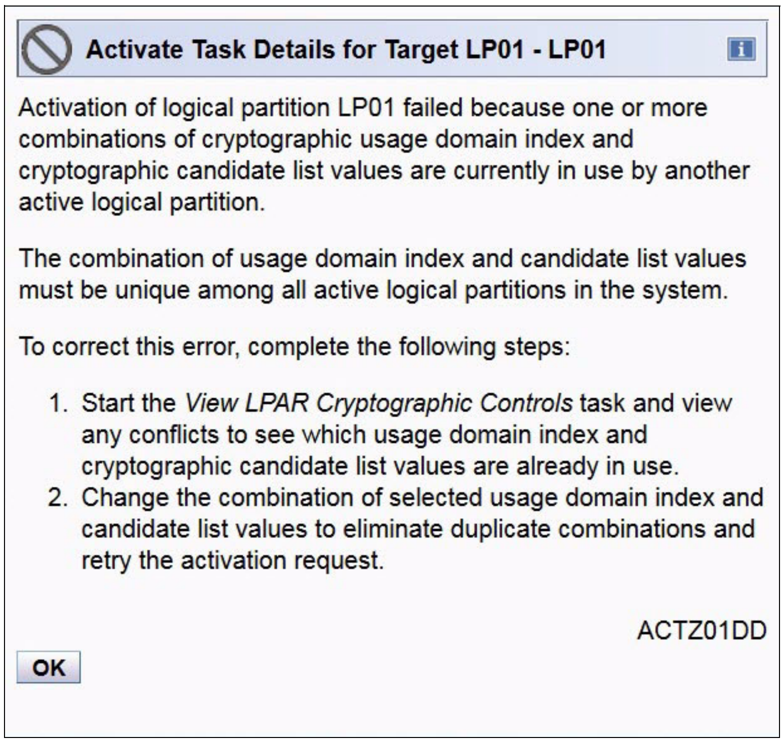

Figure 10-3 Activation of LPAR failed:ACTZ0212

- 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 later be configured online to the partition by selecting **Configure On/Off** from the SE. For more information, see "Configuring a Crypto Express5S online or offline on a logical partition" on page 562.

  When the partition is activated, no error condition is reported if a cryptographic coprocessor number 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 selected in the partition's online list was previously configured to an *off* state to the partition, it is automatically configured back to an *on* state when the partition is activated. The cryptographic online list is always selected from the image profile for each logical partition.

10.2.3 Cryptographic configuration using the Support Element (SE)

From the SE, you can do these tasks:

- Display PCI Cryptographic Configuration
- Display LPAR cryptographic controls (domain index and candidate/online lists for currently activate partitions)
- Reconfigure the coprocessor from/to accelerator
- Configure a cryptographic coprocessor and accelerator on/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 CPC Configuration section.

Figure 10-4 on page 551 shows the Cryptographic Management window. Use this window to obtain the installed cryptographic configuration (the association of the cryptographic number and the card serial number):

- View installed cryptographic features, with status and assigned PCHID and coprocessor numbers. Each PCI-X 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.
- Initiate the release of coprocessor numbers. Remove the relationship only when a Crypto Express5S feature is permanently removed from the CPC.

  The release option removes the relationship between a PCI cryptographic feature serial number and the assigned coprocessor numbers. Removing the relationship frees the coprocessor numbers, making them available to be assigned to a new feature serial number.

> **Important:** The coprocessor numbers are assigned to the feature serial number, *not* to the installed location. If a feature is removed from one location to be reinstalled in another, the coprocessor number assignment remains.

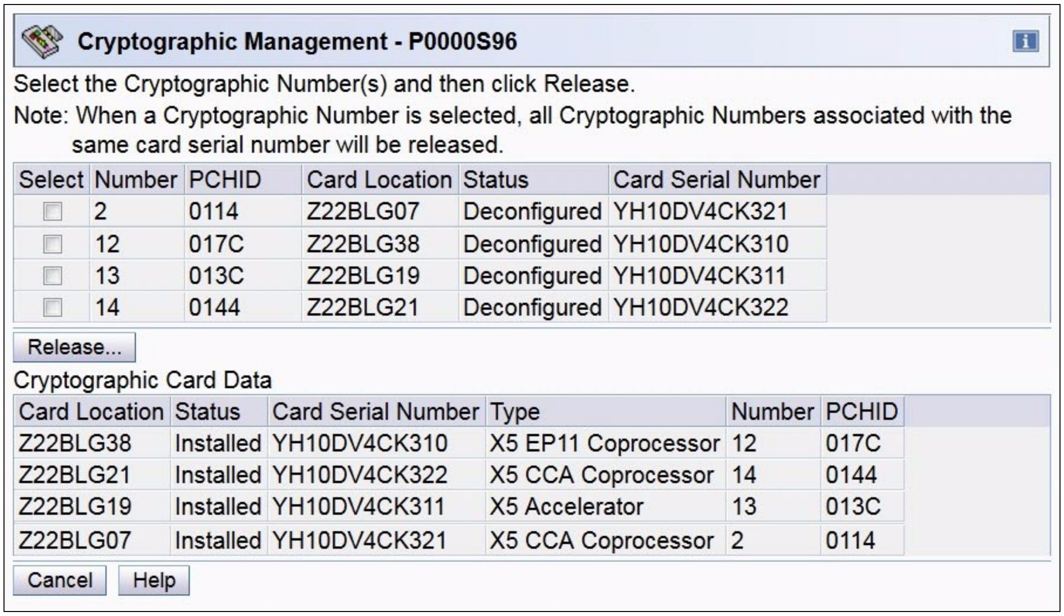

Figure 10-4 SE Cryptographic Management

View 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 CPC Operational Customization pane.

The resulting window displays the definition of Usage and Control domain indexes, and PCI Cryptographic candidate and online lists. The information is provided only for active logical partitions.

> **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 (Figure 10-5 on page 552). The usage domain index, in combination with the cryptographic number that is selected in the candidate list, must be unique across all partitions defined to the CPC. Therefore, this new tab is useful when you define or change the usage domain index for a logical partition.

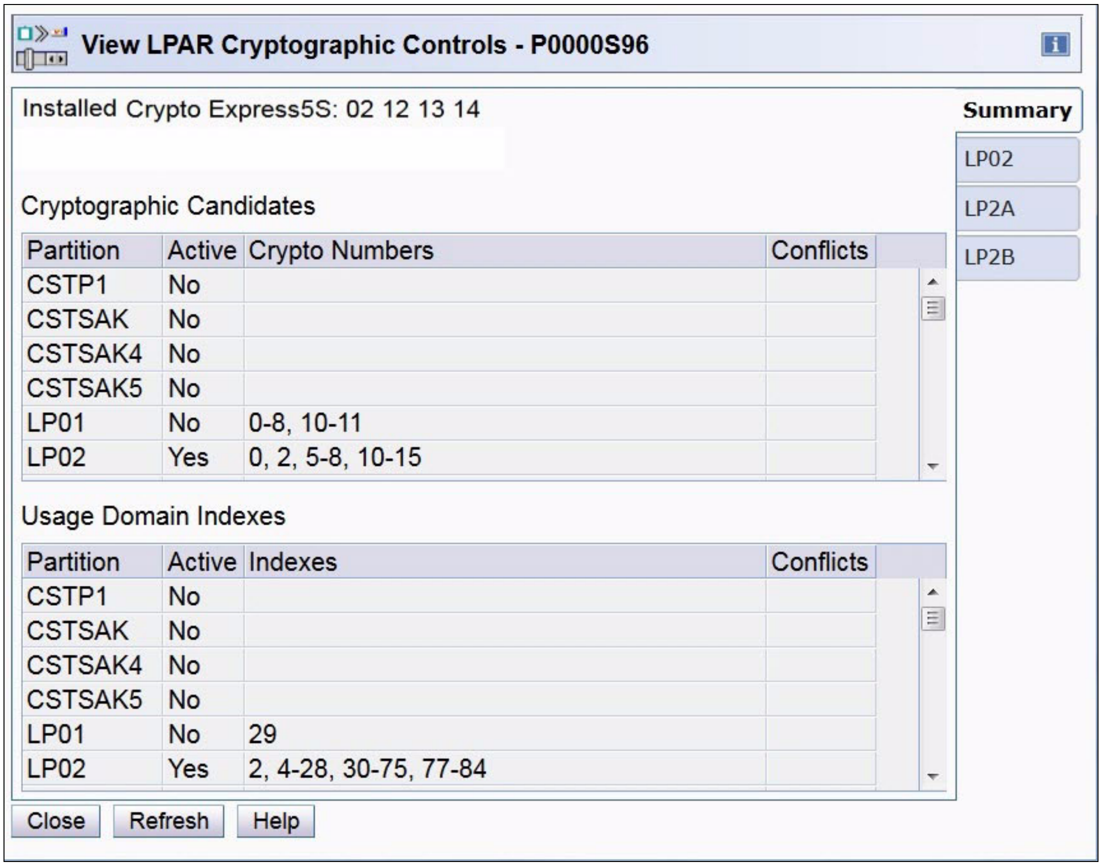

Figure 10-5 SE: View LPAR Cryptographic Controls

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 Express5S online or offline on a logical partition” on page 562.

Change LPAR Cryptographic Controls function

For each logical partition, you can define these attributes:

- Usage domain index
- Control domain index
- Cryptographic Coprocessor Candidate list
- Cryptographic Coprocessor Online list

Using the Change LPAR Cryptographic Controls function, included in the SE for the z13, you can do these tasks:

- Add a cryptographic coprocessor to a logical partition for the first time.
- Add a cryptographic coprocessor to a logical partition already using 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 panel at the Support Element. For more information about defining functions in the Image Profile, see 10.2.2, “Logical partition cryptographic definition” on page 547.

Select the Control and Usage for each domain and the cryptographic Candidate and Online for each Crypto (see Figure 10-6).

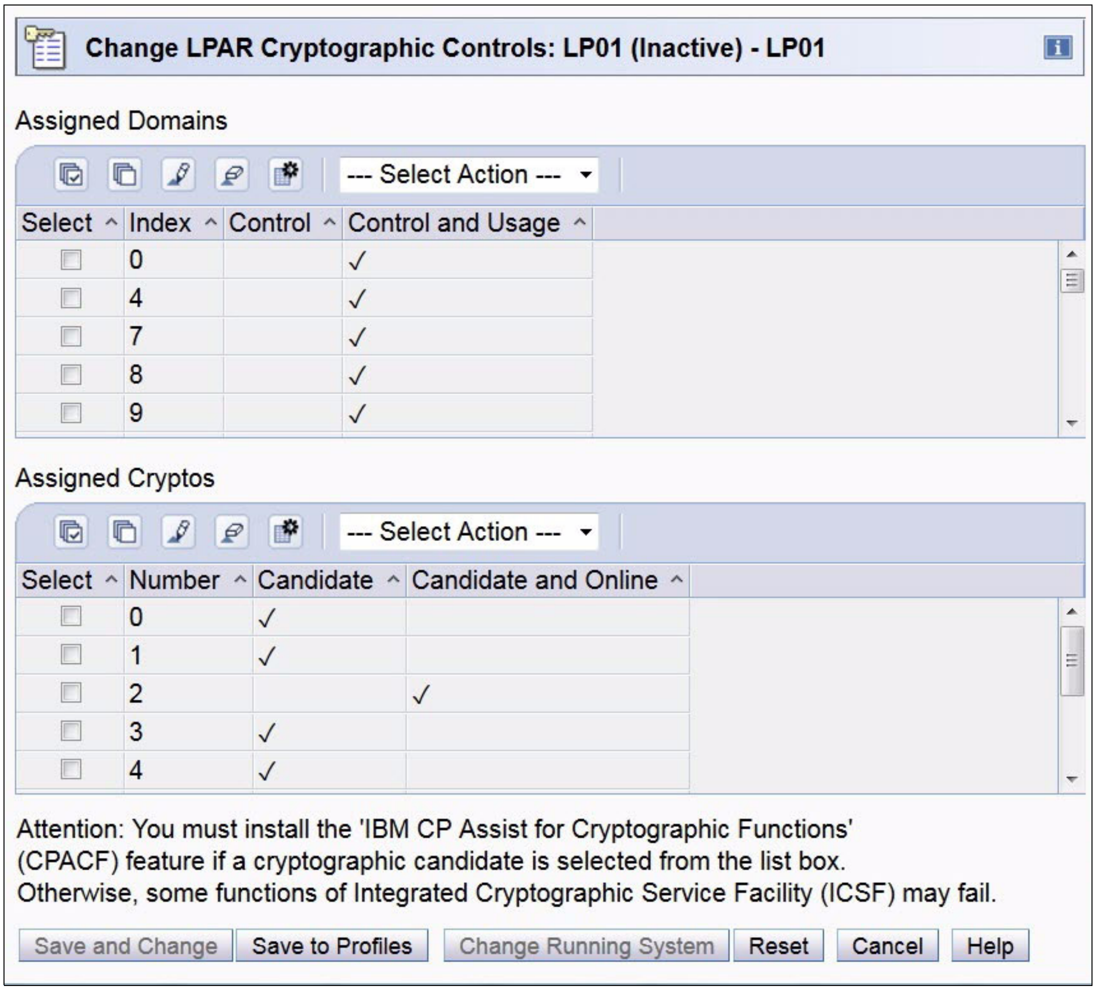

Figure 10-6 Change LPAR Cryptographic Controls: Change Running System

After selecting the appropriate boxes, you can do these tasks:

- Save these settings to the Image Profile without changing the running system.
- Change the running system without saving the definition to the Image Profile, which means after a reactivation of the partition your changes are lost.

> **Remember:** Changes to the Cryptographic Coprocessor Online List are ignored when this option is selected.

- Save the definitions to the Image Profile and activate the changes immediately to the partition.

When you add or change the control or usage domain index and cryptographic coprocessor number dynamically for a running system, a confirmation message is displayed. After processing, a status window opens and indicates the result of a dynamic addition or change of a cryptographic definition to an LPAR (see Figure 10-7).

Figure 10-7 SE: Change LPAR Cryptographic Controls

Dynamic removal of the cryptographic definition

You can remove the cryptographic definition from 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). If you try to remove the lists dynamically while they are online, the change fails and you receive the message shown in Figure 10-8.

   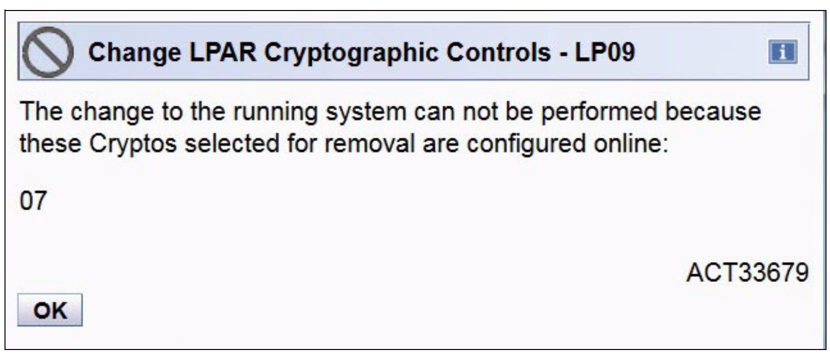

   Figure 10-8 SE: Change LPAR Cryptographic Controls: ACT33679

   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 (Figure 10-6 on page 553). 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 will also not be used for the next image activation.

2. When you remove only the definition of the cryptographic lists, the zeroize window opens (Figure 10-9).

> **Consideration:** Because you cannot see all cryptographic information, including the usage domains for other logical partitions, clear the zeroize selection by using the **Change LPAR Cryptographic Controls** task. For more information about zeroize, see “Reconfiguration of the PCI-X Adapter type” on page 556.

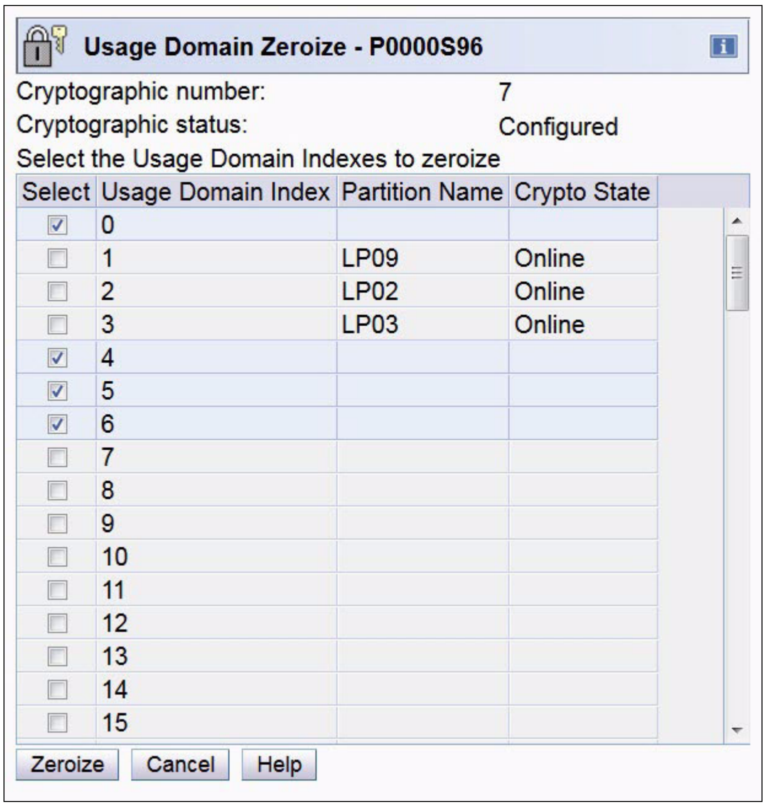

Figure 10-9 SE: Change LPAR Cryptographic Controls Zeroize

3. In the confirmation window (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.

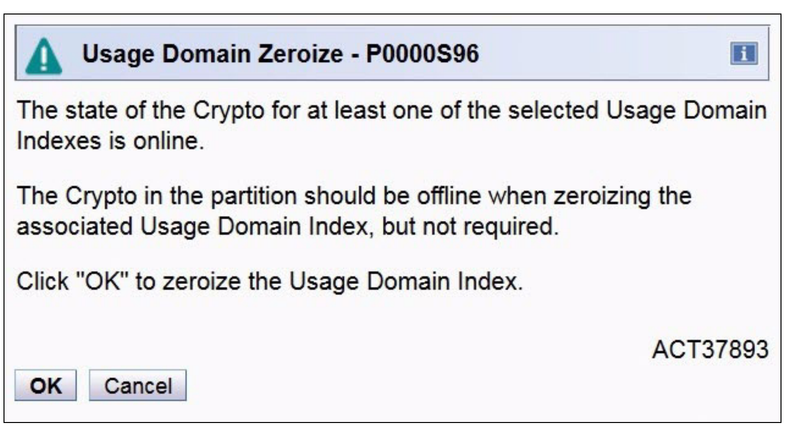

Figure 10-10 SE: Change LPAR Cryptographic Controls: ACT33680

Reconfiguration of the PCI-X Adapter type

Each PCI-X Crypto Express5S feature can be configured either as a coprocessor or as an accelerator. Each Crypto Express5S feature can be set in these configurations:

- IBM Enterprise Common Cryptographic Architecture (CCA) Coprocessor (CEX5C)
- IBM Enterprise Public Key Cryptography Standards#11 (PKCS) Coprocessor (CEX5P)
- IBM Crypto Express5S Accelerator (CEX5A)

Whether it is configured as a coprocessor or an accelerator, each PCI-X Cryptographic adapter can be shared among 85 logical partitions.

Configuring a CCA coprocessor as an accelerator

During the installation of a Crypto Express5S feature, the PCI-X Cryptographic adapter is configured by default as a CCA coprocessor. The reconfiguration is fully supported in Licensed Internal Code.

When a PCI-X adapter is configured as a CCA coprocessor, it can still run accelerator functions, albeit much more slowly than when configured as accelerator. When it is configured as an accelerator, it cannot run coprocessor functions.

When a PCI-X adapter is configured as an EP11 coprocessor, a TKE workstation is required for the management of the Crypto Express5S. For more information about configuring EP11 coprocessor, see "Configuring a CCA coprocessor as an EP11 coprocessor" on page 559.

To reconfigure the PCI-X Adapter from coprocessor to accelerator, complete these steps:

1. Select the CPC that has cryptographic coprocessor adapters you want to reconfigure, and then click the **Cryptographic Configuration** task in the CPC Configuration Group.
2. The reconfiguration is enabled only for PCI-X adapters that are Off. Therefore, be sure that the PCI-X Cryptographic adapter status for that cryptographic coprocessor channel is deconfigured (Figure 10-12 on page 557). If necessary, set the PCI-X Cryptographic adapter to Off for all partitions that have it in their candidate list. To set the PCI-X Cryptographic adapter to Off, use the procedure described in "Configuring a Crypto Express5S online or offline on a logical partition" on page 562.
3. Select the number of the cryptographic coprocessor channel (Figure 10-11) and click **Crypto Type Configuration**.

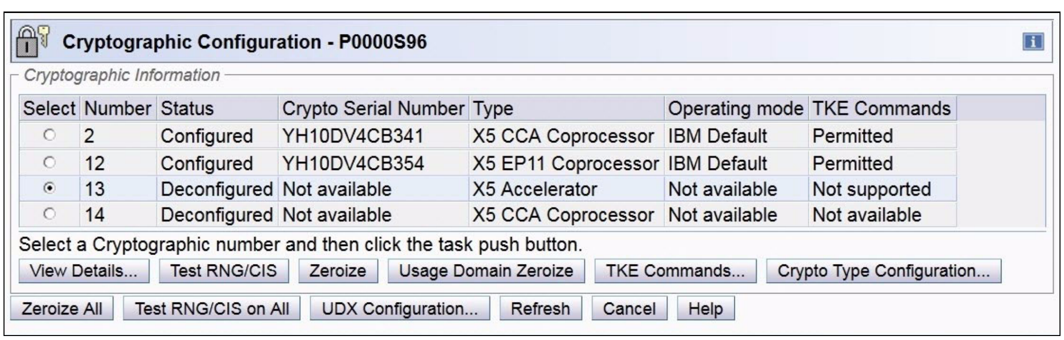

Figure 10-11 SE: Cryptographic Configuration task (Deconfigured)

4. Change the configuration for the cryptographic coprocessor adapter. The selected cryptographic coprocessor channel is configured as a coprocessor (Figure 10-12 on page 557). Select **Accelerator**.

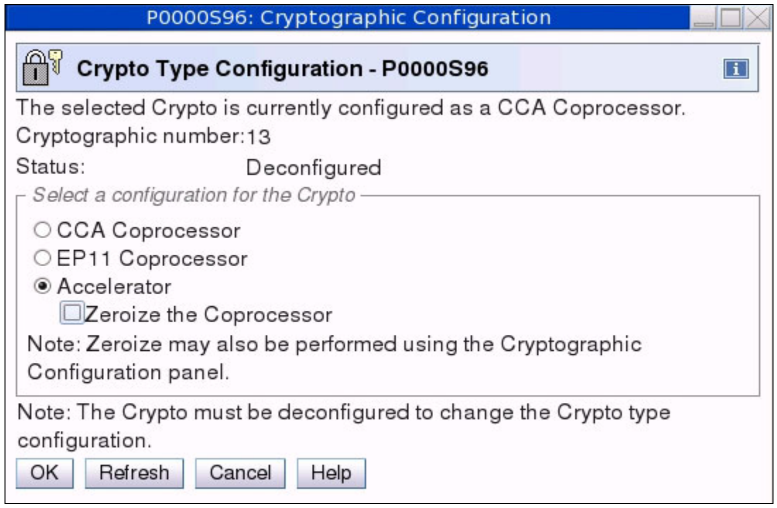

Figure 10-12 SE: Crypto Type Configuration (CCA coprocessor to Accelerator)

By selecting Accelerator, you can zeroize the selected coprocessor by also selecting **Zeroize the Coprocessor** on the Crypto Type Configuration window. However, clear the **Zeroize the Coprocessor** check box and 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 absolutely 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** (Figure 10-13).

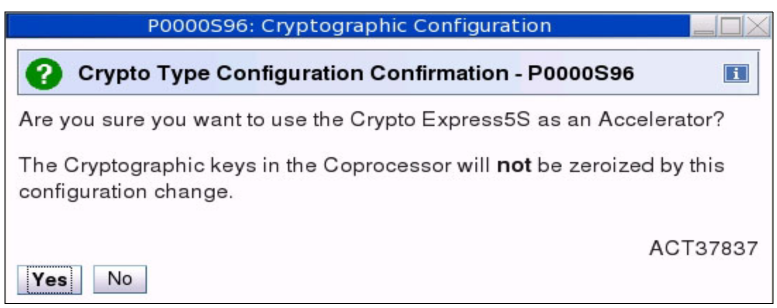

Figure 10-13 Crypto Type Configuration Confirmation for accelerator

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 and UDX Status should be `Not available` and TKE Commands `Not supported` until the cryptography is set to Online again, as described in "Configuring a Crypto Express5S online or offline on a logical partition" on page 562.

After you perform this task and go back to the Cryptographic Configuration window, the information in Figure 10-14 is displayed.

Note: UDX support is *not* available for Crypto Express5S defined as an EP11 coprocessor.

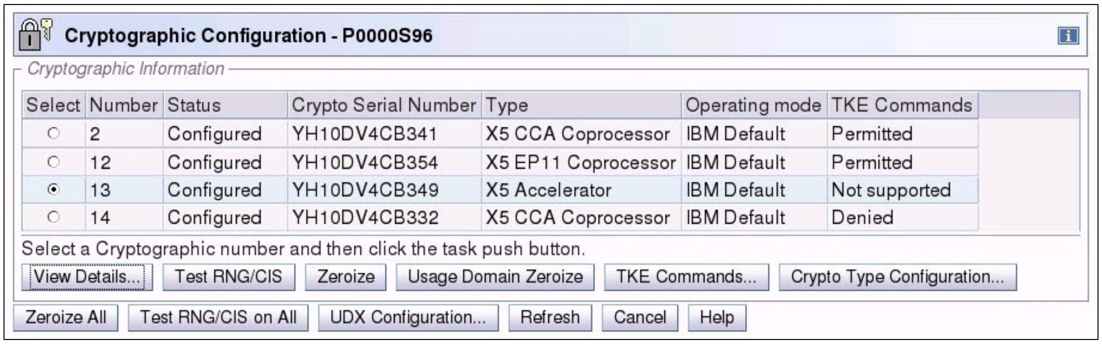

Figure 10-14 SE: Cryptographic Configuration (Accelerator online

8. Click **View Details** for detailed information (Figure 10-15).

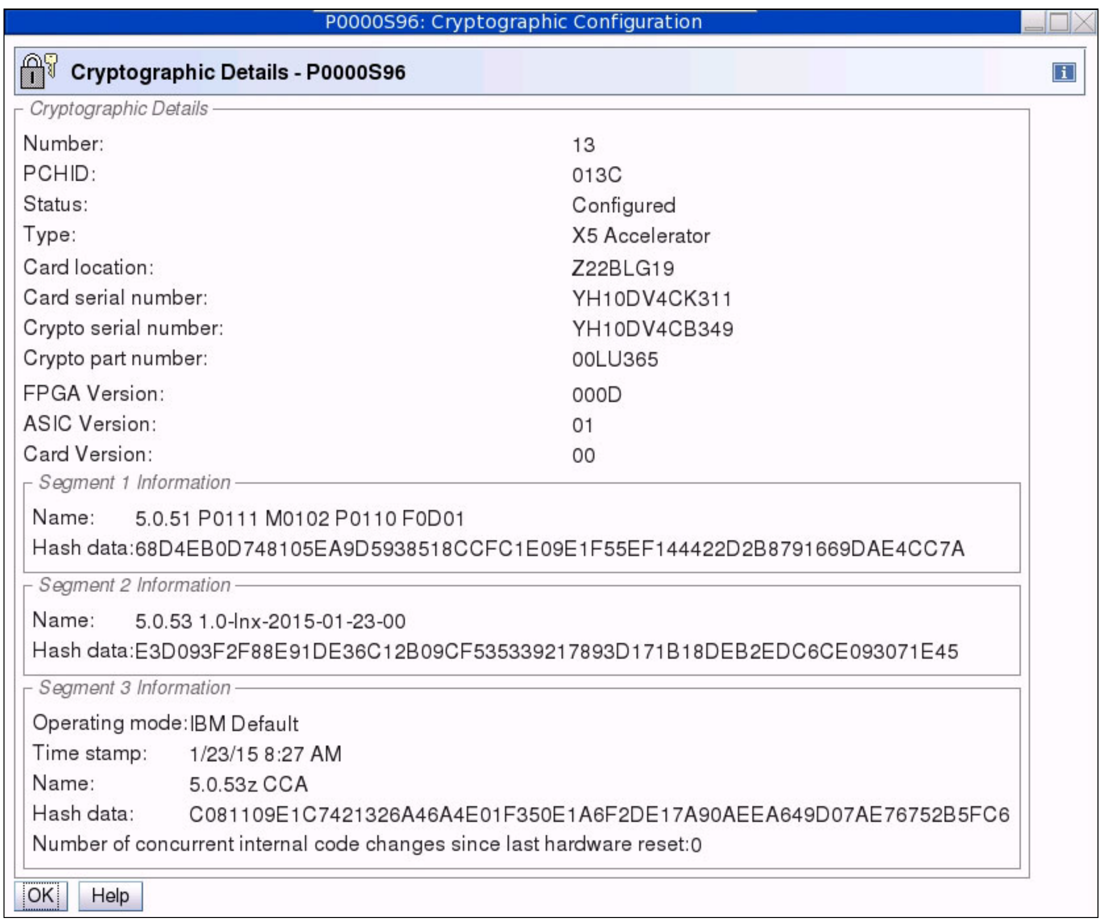

Figure 10-15 SE: Cryptographic Details (Accelerator)

The Cryptographic Type is now a Crypto Express5S 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 557.

The result of this change is shown in Figure 10-16.

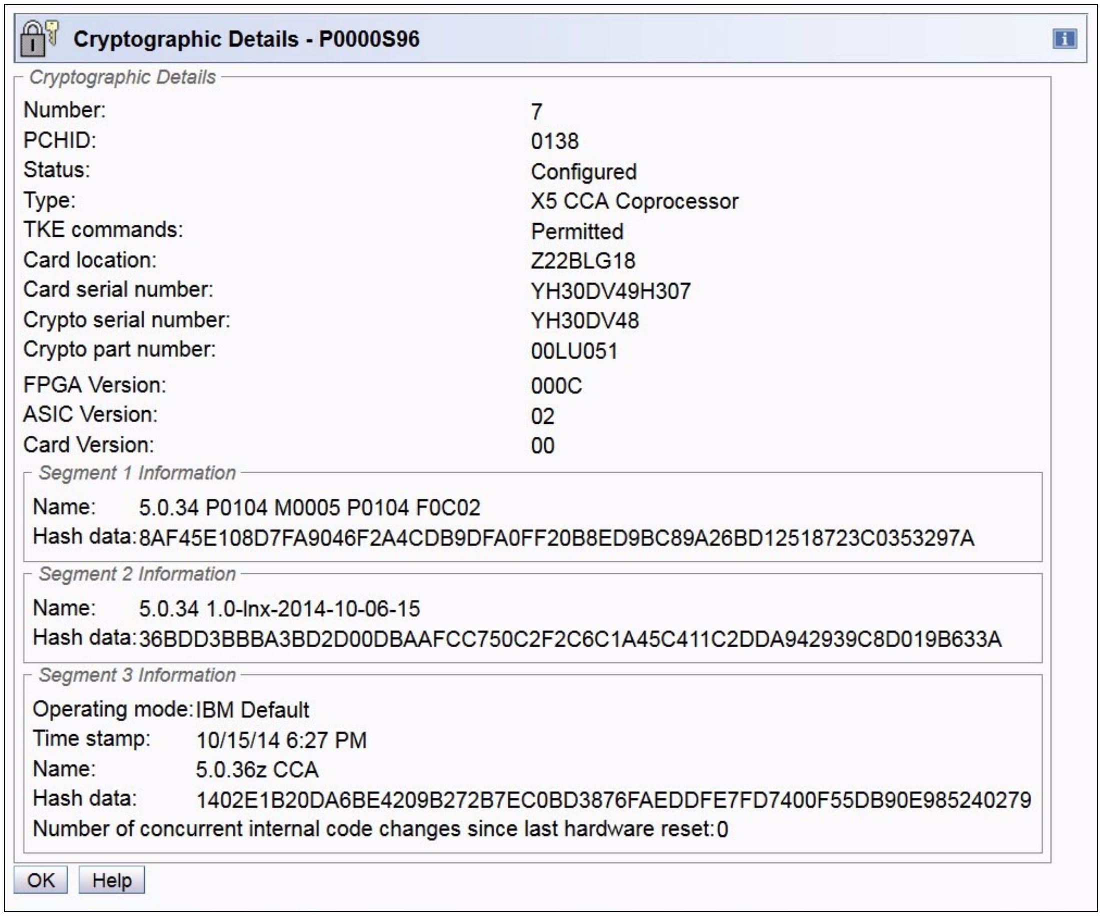

Figure 10-16 SE: Cryptographic Details (CCA Coprocessor)

Configuring a CCA coprocessor as an EP11 coprocessor

To configure a CCA coprocessor as an EP11 coprocessor, complete these steps:

1. Select the CPC that has cryptographic coprocessor adapters you want to reconfigure and click **Cryptographic Configuration** in the CPC Configuration Group.
2. The reconfiguration is enabled only for PCI-X adapters that are set to **Off**, so be sure the PCI-X Cryptographic adapter status for that cryptographic coprocessor channel is unconfigured (Figure 10-12 on page 557). If necessary, set the PCI-X Cryptographic adapter to **Off** for all partitions that have it in their candidate list. To set the PCI-X Cryptographic adapter to **Off**, use the procedure described in “Configuring a Crypto Express5S online or offline on a logical partition” on page 562.
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** (Figure 10-17), which by default, automatically selects the **Zeroize the coprocessor** option. Click **OK**.

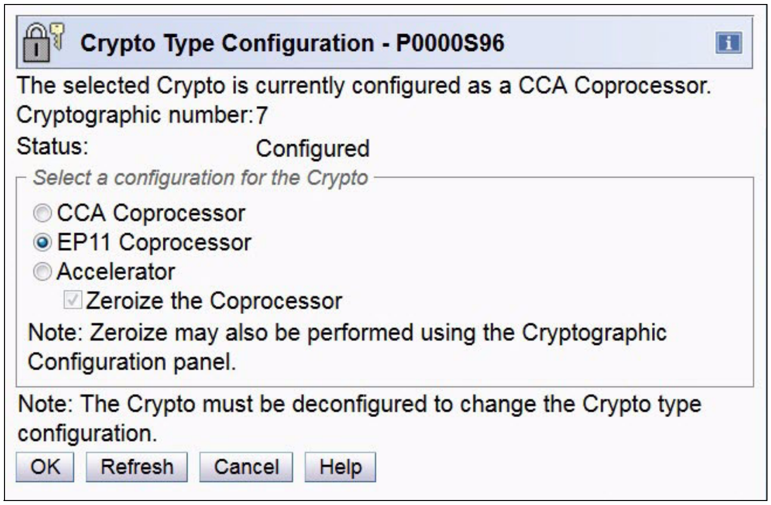

Figure 10-17 SE Crypto Type Configuration (CCA Coprocessor to EP11 Coprocessor)

5. Confirm your selection by clicking **Yes** (see Figure 10-18).

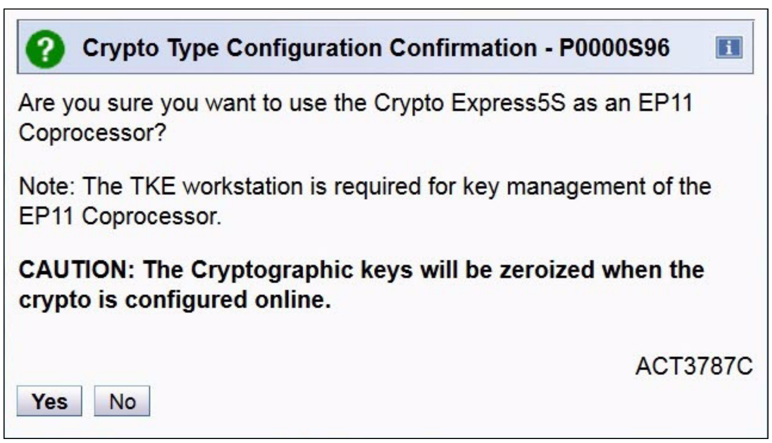

Figure 10-18 Crypto Type Configuration Confirmation for EP11 Coprocessor

6. Check that your request completed successfully. Click **OK** to return to the Crypto Type Configuration window.
7. Click **Cancel** on 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 and UDX Status should be Not available and TKE Commands Not supported until the cryptography is set Online again, as described in “Configuring a Crypto Express5S online or offline on a logical partition” on page 562.

After you complete this task and return to the Cryptographic Configuration window, the information in Figure 10-19 is displayed.

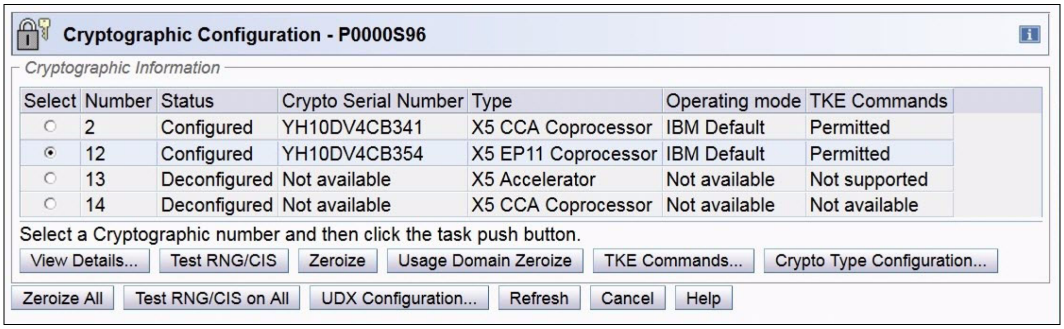

Figure 10-19 SE: Cryptographic Configuration (EP11 Coprocessor online)

8. Click **View Details**.

   Detailed information is displayed (Figure 10-20).

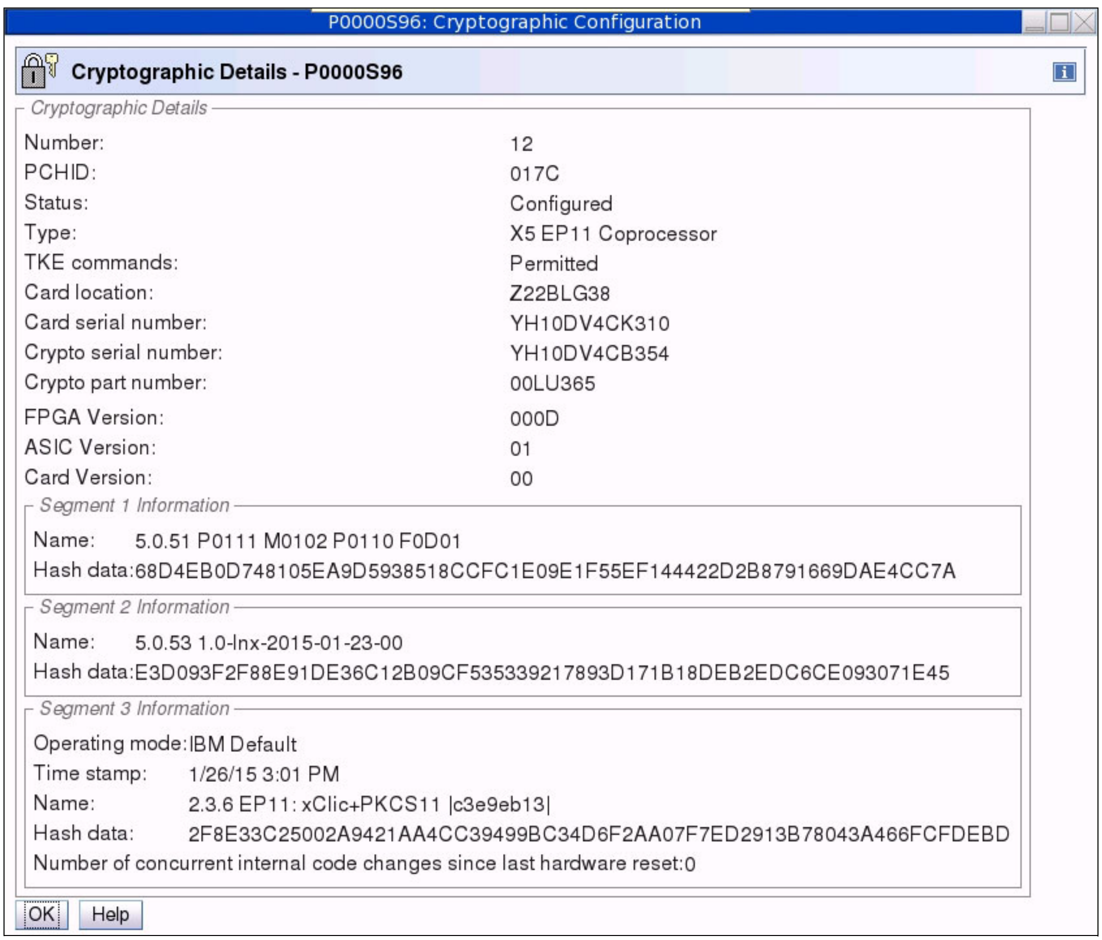

Figure 10-20 SE Cryptographic details (EP11 Coprocessor)

The Cryptographic Type is now a Crypto Express5S EP11 Coprocessor.

This 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 Express5S feature that is configured as an EP11 coprocessor, the TKE workstation is required.

Configuring a Crypto Express5S online or offline on a logical partition

For some changes to the cryptographic settings to the logical partition, you must configure the Crypto Express5S online or offline. If you can reactivate (DEACTIVATE and ACTIVATE) the image for the logical partitions whose cryptographic online lists have been updated, this dynamic operation is not needed.

Setting a Crypto Express5S to an online state

To set a Crypto Express5S online, complete the following steps:

1. From the SE, select the **Systems Management** function.
2. Select the server, click **Partitions,** and then select the target 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 two CCA Coprocessors (02 and 14), one EP11 Coprocessor (12), and one cryptographic accelerator (13) are defined to Logical Partition LP02.

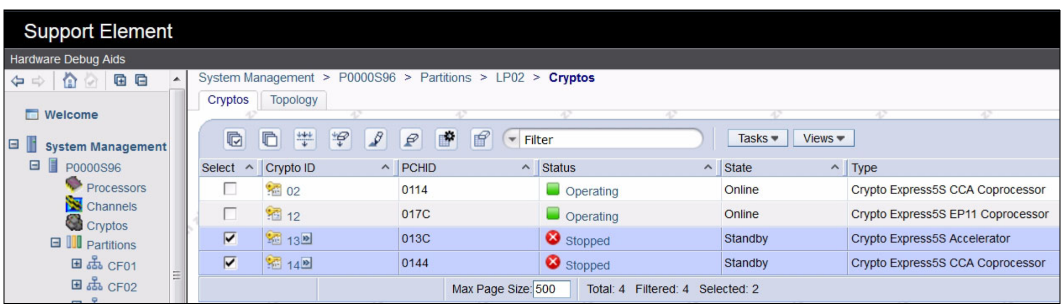

Figure 10-21 SE Crypto Service operations: Standby

5. Click **Tasks** → **Crypto Service Operations** → **Configure On/Off task** (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.

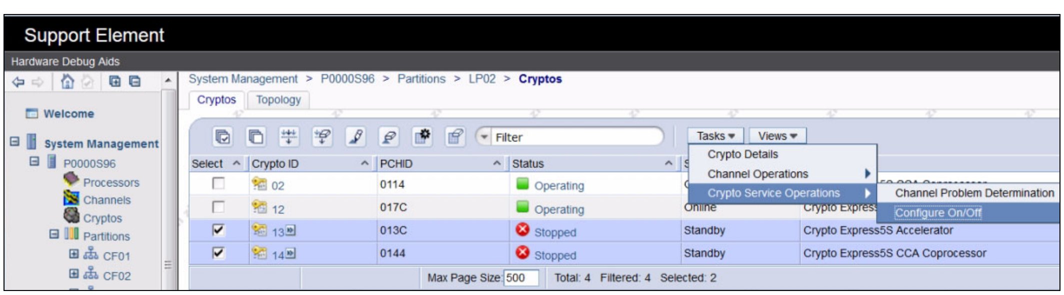

Figure 10-22 SE Crypto Service Operations: Standby

6. Select the cryptographic coprocessor channel numbers that you want, and then click **Select Action** → **Toggle** to switch from `Standby` to `Online` (Figure 10-23). If you want multiple cryptographic channels at the same time, select **Toggle All On**.

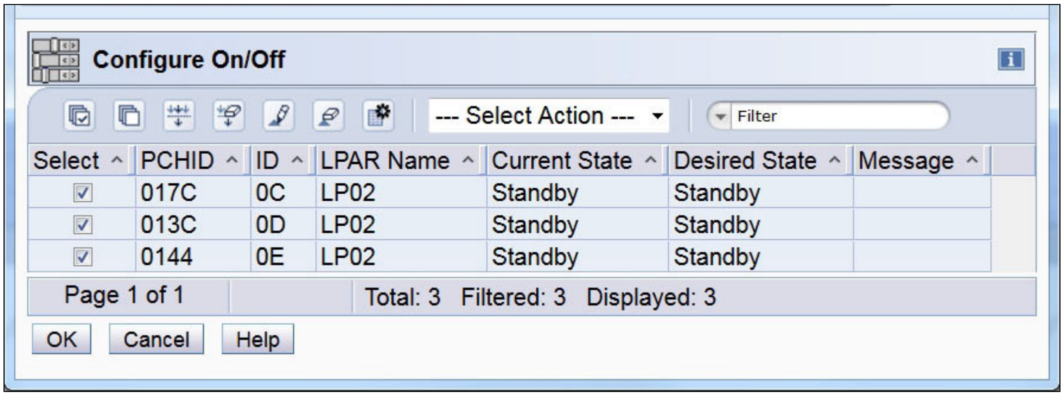

Figure 10-23 SE Config On/Off (all selected)

7. After confirming that your requested cryptographic coprocessor channels are set to the desired state of `Online`, click **OK** (Figure 10-24).

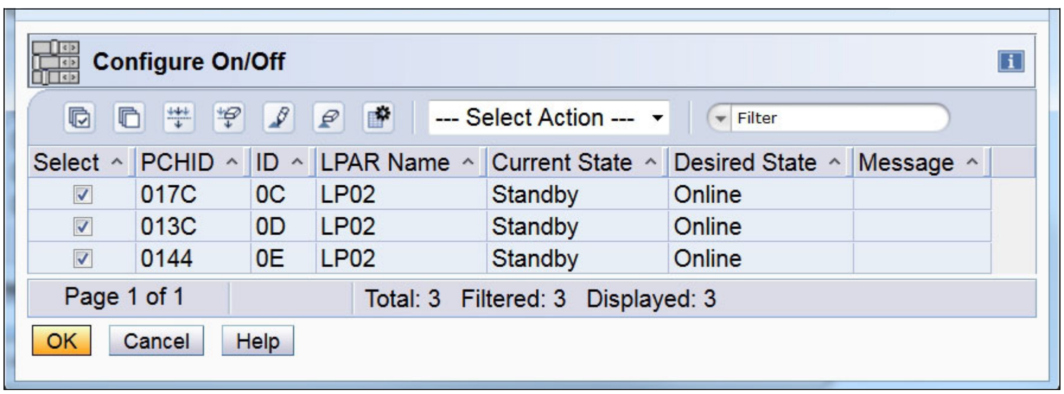

Figure 10-24 SE: Configure On/Off (Standby to Online)

8. Confirm that your request is completed (Figure 10-25). Click **OK**.

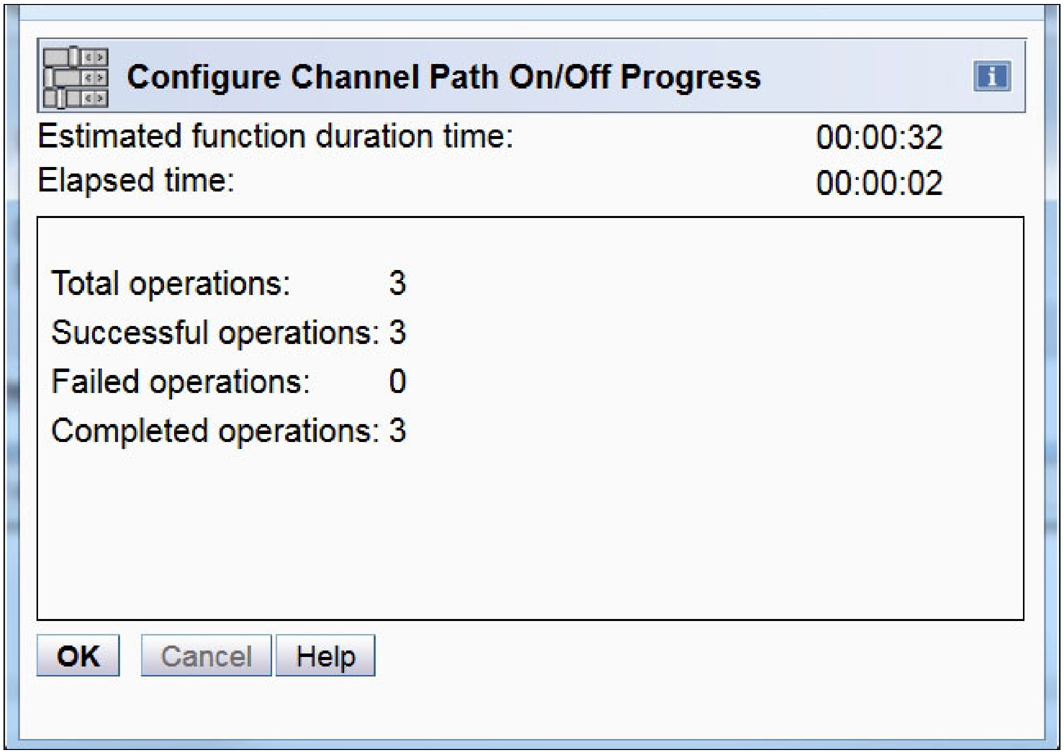

Figure 10-25 SE: Configure On/Off completed

9. After you verify that the current state of the channels changed to `Online`, click **Cancel** to return.

10. 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 (Figure 10-26).

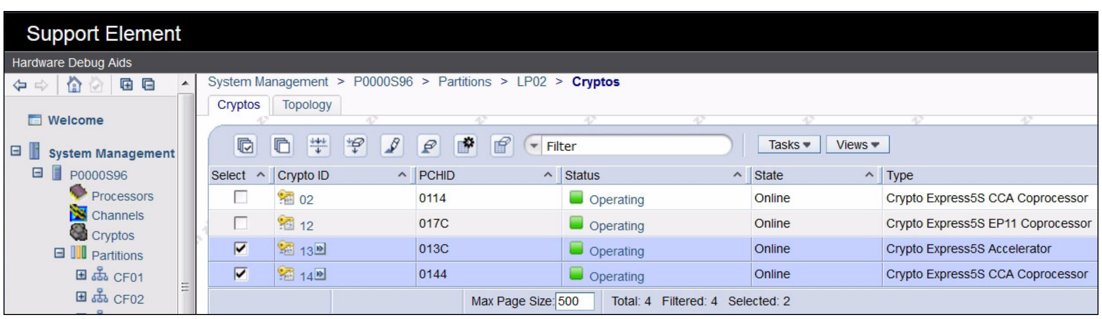

Figure 10-26 SE Crypto Service Operations: Online

Changing a cryptographic channel to standby (offline) status

To change the cryptographic channel status, complete these steps:

1. Select the logical partition whose Crypto IDs you want to change to Standby. For example, select CCA coprocessors (02), EP11 coprocessor (03), and, cryptographic accelerator (06) that are currently in an online state. Click **Tasks** → **Crypto Service Operations** → **Configure On/Off task**

2. Select the cryptographic coprocessor channel numbers that you want, and click **Select Action** → **Toggle All Standby** to switch from `Online` to `Standby` (Figure 10-27 on page 565).

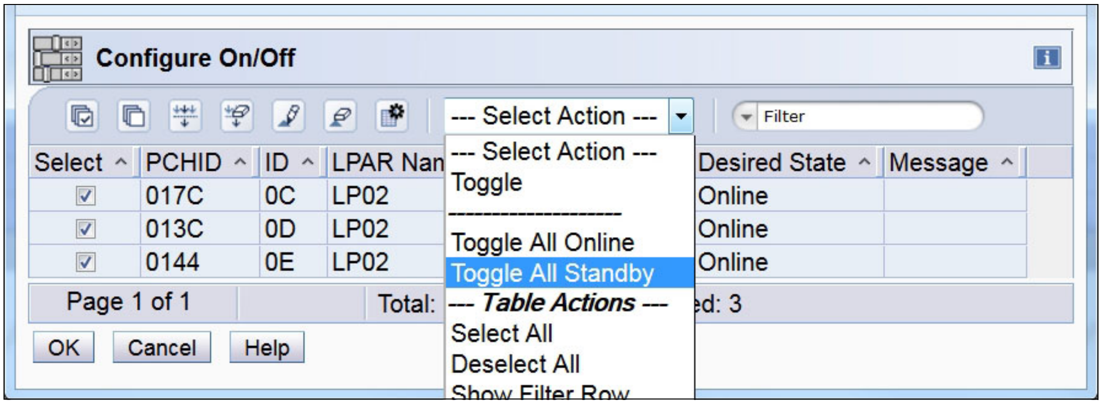

Figure 10-27 SE: Config On/Off (online - toggle all standby)

3. After you confirm that the state for your requested cryptographic channel is `Standby`, click **OK** (Figure 10-28).

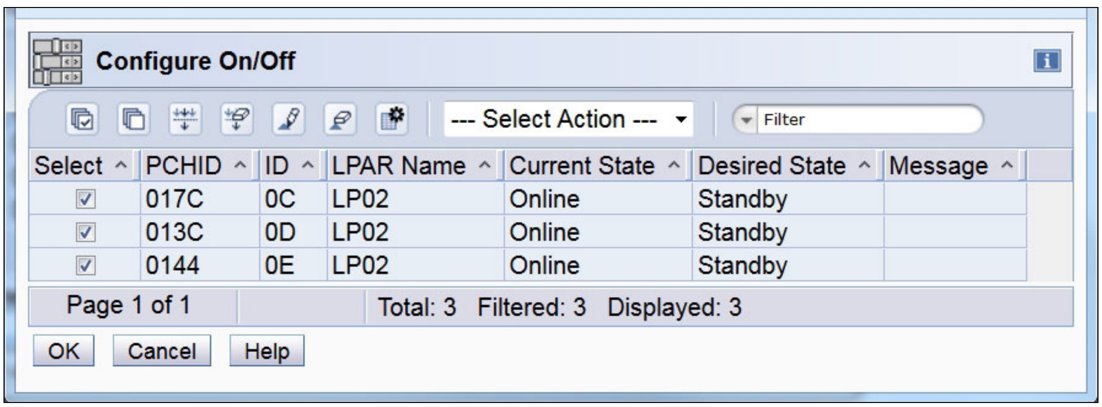

Figure 10-28 SE: Config On/Off (Online to Standby)

4. Confirm that your request is completed (Figure 10-29). Click **OK**.

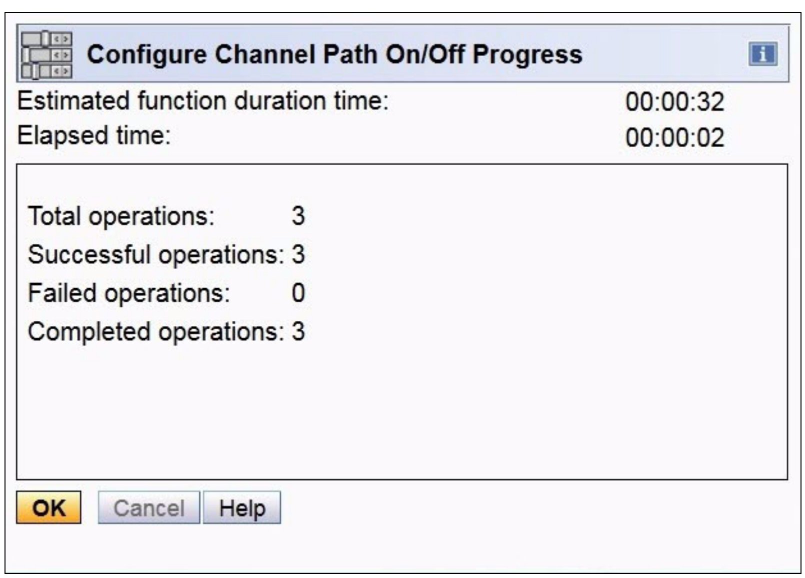

Figure 10-29 SE: Config On/Off (Standby completed)

10.3 Activating and deactivating cryptographic coprocessors by using ICSF

Integrated Cryptographic Service Facility (ICSF) provides an Interactive System Productivity Facility (ISPF) Coprocessor Management panel where you can display or change the status (Active or Deactivate) of cryptographic coprocessors. This action only affects the coprocessor status of ICSF, and has no effect on the Online/Standby hardware status that is displayed on the z13 Support Element.

From the ICSF main panel (Figure 10-30), select option **1** to open the ICSF Coprocessor Management panel.

```
HCR77B0  -------------- Integrated Cryptographic Service Facility -------------
OPTION ===>
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, 2015.
     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 10-30 Integrated Cryptographic Support Facility main panel

Cryptographic coprocessors that are currently configured on the partition are listed in the ICSF Coprocessor Management panel (Figure 10-31).

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 4 of 4
COMMAND ===>                                                  SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R and S. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS               AES   DES   ECC   RSA   P11
  -------   --------  -------------------- ---   ---   ---   ---   ---
.   5C02    DV4CB341  Active                A     A     A     A
.   5P12    DV4CB354  Active                                        A
.   5A13    N/A       Active
.   5C14    DV4CB332  Offline
******************************* Bottom of data ********************************
```

Figure 10-31 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 (Figure 10-32).

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 4 of 4
COMMAND ===>                                                  SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R and S. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS                AES   DES   ECC   RSA   P11
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   5C02    DV4CB341  Active                 A     A     A     A
.   5P12    DV4CB354  Active                                         A
.   5A13    N/A       Active
.   5C14    DV4CB332  Offline
******************************* Bottom of data ********************************
```

Figure 10-32 ICSF Coprocessor Management (Candidate only - Standby)

A cryptographic coprocessor becomes visible to ICSF Coprocessor Management when the coprocessor number is part of the partition candidate list and the coprocessor is first brought online to the partition in either of these ways:

- At the time the partition is activated, if the coprocessor is installed and the coprocessor number is part of the partition Online list.
- When the coprocessor is first configured online to the partition by using the Config On/Off task from the SE Workplace.

In the list (Figure 10-33), enter `A` or `D` to switch a coprocessor status to `Active` or `Deactivated`.

```
------------------------- ICSF Coprocessor Management -------- Row 1 to 4 of 4
COMMAND ===>                                                  SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R and S. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS                AES   DES   ECC   RSA   P11
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   5C02    DV4CB341  Active                 A     A     A     A
.   5P12    DV4CB354  Active                                         A
.   5A13    N/A       Active
d   5C14    DV4CB332  Active                 A     A     A     A
******************************* Bottom of data ********************************
```

Figure 10-33 ICSF Coprocessor Management (Online)

When a coprocessor is deactivated through ICSF (Figure 10-34 on page 568), 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 4 of 4
COMMAND ===>                                                  SCROLL ===> PAGE

 Select the cryptographic features to be processed and press ENTER.
 Action characters are: A, D, E, K, R and S. See the help panel for details.

  CRYPTO    SERIAL
  FEATURE   NUMBER    STATUS                AES   DES   ECC   RSA   P11
  -------   --------  --------------------  ---   ---   ---   ---   ---
.   5C02    DV4CB341  Active                 A     A     A     A
.   5P12    DV4CB354  Active                                         A
.   5A13    N/A       Active
.   5C14    DV4CB332  Deactivated
******************************* Bottom of data ********************************
```

Figure 10-34 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 z13 Support Element.

Figure 10-35 shows ICSF Coprocessor Hardware Status panel.

```
---------------------- ICSF - Coprocessor Hardware Status ---------------------
COMMAND ===>                                                      SCROLL ===>
                                                                 CRYPTO DOMAIN: 1

REGISTER STATUS                  COPROCESSOR 5C02
                                                                       More:     +
 Crypto Serial Number          : DV4CB341
 Status                        : ACTIVE
 AES Master Key
   New Master Key register     : EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY
     Verification pattern      :
   Current Master Key register : VALID
     Verification pattern      : 2058C870E9D3194F
 DES Master Key
   New Master Key register     : EMPTY
     Verification pattern      :
     Hash pattern              :
                               :
                               :
   Old Master Key register     : EMPTY
     Verification pattern      :
     Hash pattern              :
                               :
   Current Master Key register : VALID
     Verification pattern      : CA6B408A02371B1D
     Hash pattern              : DF3A50AE35466123
                               : 96EF557E8BD074C1
 ECC Master Key
   New Master Key register     : EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY
     Verification pattern      :
   Old Master Key register     : EMPTY
     Verification pattern      :
   Current Master Key register : VALID
     Verification pattern      : 78D81AC6C9610A2C
 RSA Master Key
   New Master Key register     : EMPTY
     Verification pattern      :
                               :
   Old Master Key register     : EMPTY
     Verification pattern      :
                               :
   Current Master Key register : VALID
     Verification pattern      : E83F158521FEEA23
                               : 986CC9483DAFD711

Press ENTER to refresh the hardware status display.
Press END   to exit to the previous menu.
```

Figure 10-35 ICSF Coprocessor Hardware Status

Help information from ICSF Coprocessor Management (Figure 10-36 and Figure 10-37 on page 571), 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
  2C        Crypto Express2 Coprocessor              a, d, e, k, r, s
  2A        Crypto Express2 Accelerator              a, d,
  3C        Crypto Express3 Coprocessor              a, d, e, k, r, s
  3A        Crypto Express3 Accelerator              a, d
  4A        Crypto Express4 Accelerator              a, d
  4C        Crypto Express4 CCA coprocessor          a, d, e, k, r, s
  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
  5P        Crypto Express5 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.
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.
```

Figure 10-36 Help for Coprocessor Management (part 1 of 2)

```
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.
                           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.
  - 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 10-37 Help for Coprocessor Management (part 2 of 2)

11

Logical partitions

This chapter describes how to implement and use reserved logical partitions on the IBM z13.

This chapter includes the following topics:

- Introduction to reserved logical partitions
- Naming and activating a reserved logical partition
- Changing a logical partition to reserved
- Renaming a logical partition

11.1 Introduction to reserved logical partitions

A logical partition can be named, given resources, and activated when needed. A logical partition can also be deactivated and unnamed, meaning it becomes a reserved logical partition. In addition, a logical partition can have its name changed. All of these changes can be done without the need for a power-on reset (POR).

Various reasons exist for why you might want to use reserved logical partitions. As the capacity of one server increases, an associated tendency is for the number of logical partitions on the server to increase. In this case, you might need to add a z/OS image to remove redundant z/OS images, or consolidate logical partitions from multiple servers. All of these tasks can be run without disrupting the operations of the logical partitions. The use of reserved logical partitions also includes logical partitions that are used as coupling facility images or Linux partitions.

The use of reserved logical partitions, with concurrent book, memory, and channel upgrades, means that the need for a POR of the z13 is kept to a minimum. This in turn means that unnecessary downtime on active logical partitions is also reduced.

During the initial configuration of a processor, you must plan for several hardware configuration definitions, including the logical partitions, channels, channel subsystems (CSSs), and subchannel sets. The reserved logical partitions also require extra space in the HSA. However, you are no longer required to plan for hardware system area (HSA) growth because of extra definitions while planning for the z13.

Important: On z13, a fixed amount (96 GB) is reserved for HSA storage, and is separated from client purchased memory. When you define a z13 in Hardware Configuration Definition (HCD), the maximum numbers of logical partitions, channel subsystems, and subchannel sets are defined. When the associated IOCDS is activated, HSA space is allocated for this configuration. Therefore, you do not need to reserve space for future partitions in the input/output configuration data set (IOCDS) or do a POR even if a new logical partition is needed.

The ability to use reserved logical partitions depends on hardware and software requirements being in place before this function can be used. In addition, any independent software vendor (ISV) software that relies on the multiple image facility ID (MIFID) or partition ID might need to be updated.

11.2 Naming and activating a reserved logical partition

To activate a reserved logical partition, an HCD configuration change is required to assign a name to the logical partition. In addition, the image profile for the new logical partition must be created or updated.

The assumption is that the z13 has sufficient resources for the new logical partition, which includes memory, channels, and CPs. Use the following steps, which are described in the next sections, to name and activate a reserved logical partition and customizing the profile:

- Naming a reserved logical partition
- Activating the configuration on the z13
- Customizing the Image Profile

11.2.1 Naming a reserved logical partition

To name the reserved logical partition, complete the following steps using HCD:

1. From the HCD main menu, select option **1. Define, modify, or view configuration data**.
2. Select option **3. Processors** to display the processor list**.**
3. Select the processor that has the reserved logical partition you want to name by using the **work with partitions** (p) option.
4. Select the CSS from the channel subsystem list that contains the reserved logical partition that you want to name. This action displays the partition list for the selected CSS by using the **work with partitions** (p) option.
5. From the Partition List, name the reserved logical partition by using the **Change** (c) option. The example in Figure 11-1 shows that partition Number 4 in Channel Subsystem ID 1 is reserved and that its Partition Name is displayed as an asterisk (*). HCD sometimes displays the MIFID as a partition number or just a number.

```
----------------------------- Partition List -----------------------------
   Goto  Backup  Query  Help
 ------------------------------------------------------------------------
                                                                Row 1 of 15
 Command ===> ________________________________________ Scroll ===> PAGE

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

 Processor ID  . . . . : SCZP501   Sphinx
 Configuration mode  . : LPAR
 Channel Subsystem ID  : 1

 / Partition Name   Number Usage + Description
 _ A1B              B      CF/OS   CHPID holder
 _ A1C              C      CF      WTSCPLX8 CF8C
 _ A11              1      OS      WTSCPLX8 SC81
 _ *                2      CF/OS   ________________________________
 _ *                3      CF/OS   ________________________________
 _ *                4      CF/OS   ________________________________
 _ *                5      CF/OS   ________________________________
```

Figure 11-1 HCD: Changing a reserved logical partition

Figure 11-2 shows that the logical partition is given the new name A14. The partition number (MIF image ID) cannot change from what it was set to in the initial partition definition when this partition was created. This value already is set when the processor was power-on reset.

```
-------------- Change Partition Definition ---------------


 Specify or revise the following values.

 Partition name . . . A14
 Partition number . . 4      (same as MIF image ID)
 Partition usage  . . CF/OS  +

 Description  . . . . Dynamically Added LPAR SCZP501
```

Figure 11-2 HCD: Naming a reserved logical partition

6. Change Partition usage, if necessary.
7. Assign all channels that are required by the logical partition to access the devices it needs by using the **Include in access list** (a) option or by entering `YES` in the Access List column (Figure 11-3). All devices on the selected CHPIDs are available to the newly named logical partition except those with an explicit device candidate list.

```
------------- Update CHPID Access and Candidate Lists -------------
                                                      Row 34 of 107
 Command ===> _________________________________ Scroll ===> PAGE

 Select channel paths to include or exclude, then press Enter.

 Partition name . . : A14       Dynamically Added LPAR SCZP501

                        ---Partition Included In---
 / CHPID  Type   Mode   Access List  Candidate List
 _ 4D     FC     SPAN   No           No
 _ 4E     FC     SPAN   No           No
 _ 4F     FC     SPAN   No           No
 a 50     FC     SPAN   No           No
 a 51     FC     SPAN   No           No
 a 52     FC     SPAN   No           No
 a 53     FC     SPAN   No           No
 _ 54     FC     SPAN   No           No
 _ 55     FC     SPAN   No           No
 _ 56     FC     SPAN   No           No
 _ 57     FC     SPAN   No           No
```

Figure 11-3 HCD: Add/Assign CHPID to new named logical partition

You might need to take more steps to define devices such as channel-to-channel connections (CTCs) and Coupling links, including Coupling over InfiniBand (CIB). An operating system configuration might have to be created if you use a different one from those already defined. Remember to ensure that devices such as disk subsystems have sufficient logical channel connections to support the new logical partition I/O activity.

8. If you have channels that you do not want to connect to the newly named logical partition, those channels must have their candidate lists updated to exclude the newly named logical partition. If some devices already have candidate lists, the new logical partition is not included in that candidate list.

Your partition is now named and has all the necessary resources defined from an HCD perspective (Figure 11-4). At this point, a production input/output definition file (IODF) can be built.

```
----------------------------- Partition List -----------------------------
   Goto  Backup  Query  Help
 ------------------------------------------------------------------------
                                                               Row 1 of 15
 Command ===> ________________________________________ Scroll ===> PAGE

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

 Processor ID  . . . . : SCZP501   Sphinx
 Configuration mode  . : LPAR
 Channel Subsystem ID  : 1

 / Partition Name   Number Usage + Description
 _ A1B              B      CF/OS   CHPID holder
 _ A1C              C      CF      WTSCPLX8 CF8C
 _ A11              1      OS      WTSCPLX8 SC81
 _ A14              4      CF/OS   Dynamically Added LPAR SCZP501
 _ *                2      CF/OS   ________________________________
 _ *                3      CF/OS   ________________________________
 _ *                5      CF/OS   ________________________________
```

Figure 11-4 HCD: Partition List with new named logical partition

11.2.2 Activating the configuration on the z13

This section explains how to dynamically activate the production IODF to the processor and the systems that run on it. For more information, see the *Hardware Configuration Definition User's Guide*, SC33-7988.

Checking the definition for IODF dynamic change

To check the definition, complete the following steps:

1. Before dynamically activating the production IODF that contains the definition changes of the newly named logical partition, verify that the processor is enabled for dynamic activation. From the Dynamic section in the activation profile of the CPC on the HMC application, verify that the **Allow dynamic changes to the channel subsystem I/O definition** option is selected.

   This option *must* have already been enabled before the last power-on reset of the z13 (Figure 11-5). If not, a power-on reset is needed to change the reserved logical partition to a named logical partition.

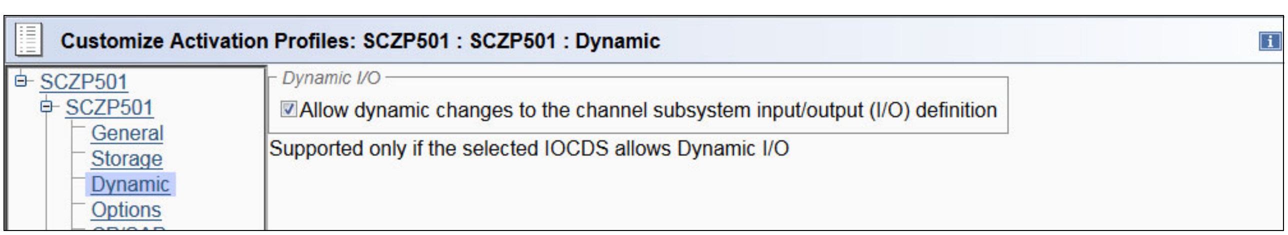

Figure 11-5 The HMC: Dynamic section in reset profile

2. Check that the z/OS system supports the command to dynamically activate the new configuration. This command is issued through HCD or as a z/OS **activate** command. When that command is issued in the system running on z13 for which z/OS V1.12 (lifecycle extension required) or later is required, the software requirement will already be satisfied.

> **The activate command:** Before you issue the **activate** command, be sure to verify that full software and hardware activation can be done. From a system console, issue the **d ios,config(all)** z/OS command. A panel similar to the one in Figure 11-6 is displayed. This panel shows that the source IODF and the active IODF are the same (SYS6.IODF42). Therefore, a full activation can be done.

```
D IOS,CONFIG(ALL)
IOS506I 09.12.57 I/O CONFIG DATA 219
ACTIVE IODF DATA SET = SYS6.IODF42
CONFIGURATION ID = L06RMVS1       EDT ID = 01
TOKEN:  PROCESSOR DATE     TIME     DESCRIPTION
 SOURCE: SCZP501  14-11-07 10:15:41 SYS6     IODF42
ACTIVE CSS:  0    SUBCHANNEL SETS CONFIGURED: 0, 1, 2, 3
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE
LOCAL SYSTEM NAME (LSYSTEM): SCZP501
HARDWARE SYSTEM AREA AVAILABLE FOR CONFIGURATION CHANGES
PHYSICAL CONTROL UNITS              8093
CSS  0 - LOGICAL CONTROL UNITS      3992
 SS  0   SUBCHANNELS               54701
 SS  1   SUBCHANNELS               58862
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
CSS  1 - LOGICAL CONTROL UNITS      3998
 SS  0   SUBCHANNELS               55001
 SS  1   SUBCHANNELS               58862
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
CSS  2 - LOGICAL CONTROL UNITS      4088
 SS  0   SUBCHANNELS               65280
 SS  1   SUBCHANNELS               65535
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
CSS  3 - LOGICAL CONTROL UNITS      4088
 SS  0   SUBCHANNELS               65280
 SS  1   SUBCHANNELS               65535
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
CSS  4 - LOGICAL CONTROL UNITS      4088
 SS  0   SUBCHANNELS               65280
 SS  1   SUBCHANNELS               65535
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
CSS  5 - LOGICAL CONTROL UNITS      4088
 SS  0   SUBCHANNELS               65280
 SS  1   SUBCHANNELS               65535
 SS  2   SUBCHANNELS               65535
 SS  3   SUBCHANNELS               65535
ELIGIBLE DEVICE TABLE LATCH COUNTS
        0 OUTSTANDING BINDS ON PRIMARY EDT
```

Figure 11-6 Display IOS,CONFIG output

3. You can also use the HCD **Activate or select Process configuration data**, then the **View Active Configuration** option to verify that a full activation can be run (Figure 11-7).

```
------------------------- View Active Configuration -------------------------


 Currently active IODF  . . : SYS6.IODF42
    Creation date . . . . . : 14-11-07
    Volume serial number  . : IODFPK

 Configuration ID . . . . . : L06RMVS1     Sysplex systems
 EDT ID . . . . . . . . . . : 01

 HSA token  . . . . . . . . : SCZP501  14-11-07 10:15:41 SYS6     IODF42

 Activation scope:
 Hardware changes allowed . : Yes
 Software changes allowed . : Yes

 ENTER to view details on the activation scope.
```

Figure 11-7 HCD: View Active Configuration

Using this option displays the CBDA781I message (Figure 11-8).

```
------------------------------- Message List --------------------------------
   Save  Query  Help
 ---------------------------------------------------------------------------
                                                                  Row 1 of 2
 Command ===> ___________________________________________ Scroll ===> PAGE

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

 / Sev Msg. ID  Message Text
 _ I   CBDA781I Your system configuration provides full dynamic
 #              reconfiguration capability.
 **************************** Bottom of data ******************************
```

Figure 11-8 HCD: Display message for View Active Configuration

Activating a new IODF

After verification of the current definition, the new configuration can be activated dynamically. Watch for messages and correct any errors that occur. The dynamic activation can be run either by using the z/OS **`activate`** command or by HCD.

To activate a new IODF, complete the following steps:

1. In HCD, select option **2. Activate or process configuration data**.
2. Select option **6. Activate or verify configuration dynamically**.
3. Select option **1. Activate new hardware and software configuration**, which allows you to do a full dynamic activation (Figure 11-9).

```
--------------- Activate or Verify Configuration -------------


The currently active IODF matches the hardware I/O
configuration.  Both hardware and software definitions may be
changed. Select one of the following tasks.

1_  1.  Activate new hardware and software configuration.
    2.  Activate software configuration only. Validate
        hardware changes. Process changes to Coupling
        Facility elements.
    3.  Activate software configuration only.
    4.  *Verify active configuration against system.
    5.  *Verify target configuration against system.
    6.  Build CONFIGxx member.

* = requires TSA I/O Operations
```

Figure 11-9 HCD: Activate or Verify Configuration window

4. Verify that the new production IODF can be successfully activated. To do this, set the **Test only** option to `Yes` in HCD (Figure 11-10), or use the **`activate iodf=xx,test`** z/OS command.
5. To activate the configuration, change the **Test only** option to `No`, so that the new hardware and software configuration is activated at the same time (Figure 11-10).

```
------------- Activate New Hardware and Software Configuration -------


 Specify or revise the values for IODF activation.

 Currently active IODF . : SYS6.IODF42
    Processor ID . . . . : SCZP501      Sphinx
    Configuration ID . . : L06RMVS1     Sysplex systems
    EDT ID . . . . . . . : 01

 IODF to be activated  . : SYS6.IODFE7
    Processor ID . . . . . SCZP501   +
    Configuration ID . . . L06RMVS1  +     EDT ID . . . 01  +

 Test only  . . . . . . . . . . . . . . . . . . . . . No   (Yes or No)
 Allow hardware deletes (FORCE, FORCE=DEVICE) . . . . No   (Yes or No)
 Delete partition access to CHPIDs unconditionally
 (FORCE=CANDIDATE)  . . . . . . . . . . . . . . . . . No   (Yes or No)
 Write IOCDS  . . . . . . . . . . . . . . . . . . . . No   (Yes or No)
 Switch IOCDS for next POR  . . . . . . . . . . . . . No   (Yes or No)
```

Figure 11-10 HCD: Activate new IODF

6. Review all messages in the resulting panel (Figure 11-11), and verify that the dynamic activation completed successfully. The action performed by this HCD panel is the same as issuing the **activate iodf=xx** z/OS command.

```
------------------------------- Message List --------------------------------
   Save  Query  Help
 --------------------------------------------------------------------------
                                                                 Row 1 of 11
 Command ===> ___________________________________________ Scroll ===> PAGE

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

 / Sev Msg. ID  Message Text
 _ I   IOS500I  ACTIVATE RESULTS
 #              ACTIVATE COMPLETED SUCCESSFULLY
 _ I   CBDA888I Following devices are to be modified for processor
 #              SCZP501: 1.2100-1.2108,1.210F,1.2120-1.2128,1.212F
 _ I   CBDA882I Following channel paths are to be modified for processor
 #              SCZP501:
 #              1.00-1.02,1.07-1.09,1.0B-1.0D,1.10-1.13,1.40-1.48,1.4F-1.
 #              53,1.60-1.68,1.F3-1.F4,1.FE-1.FF
 _ I   CBDA899I Following partitions are to be added to processor
 #              SCZP501: 1.A14
 _ I   CBDA126I ACTIVATE IODF=E7 command was accepted.
 **************************** Bottom of data ******************************
```

Figure 11-11 HCD: Activate new IODF message

7. Confirm the current active IODF by using the **View active configuration** option in HCD or the **d ios,config** z/OS command. Now you can view the newly defined logical partition name in the contents of the CPC when you log in to the HMC.
8. Install a new IOCDS to the z13 and make it the active IOCDS on the server for the next POR.
9. From HCD, select option **2. Activate or process configuration data**.
10. Select option **11. Build and manage System z cluster IOCDS, IPL attributes and dynamic I/O changes** to update the new IOCDS into the HSA token slot and switch the active IOCDS. This operation enables you to use the new IOCDS, which contains the newly named logical partition, if a POR is required.

11.2.3 Customizing the Image Profile

To customize the Image Profile for the new named logical partition, complete these steps:

1. From the HMC, select the **Customize/Delete Activation Profile** for the system that contains the newly named logical partition.
2. Select the reset profile you want and click **Customize profile** (Figure 11-12).

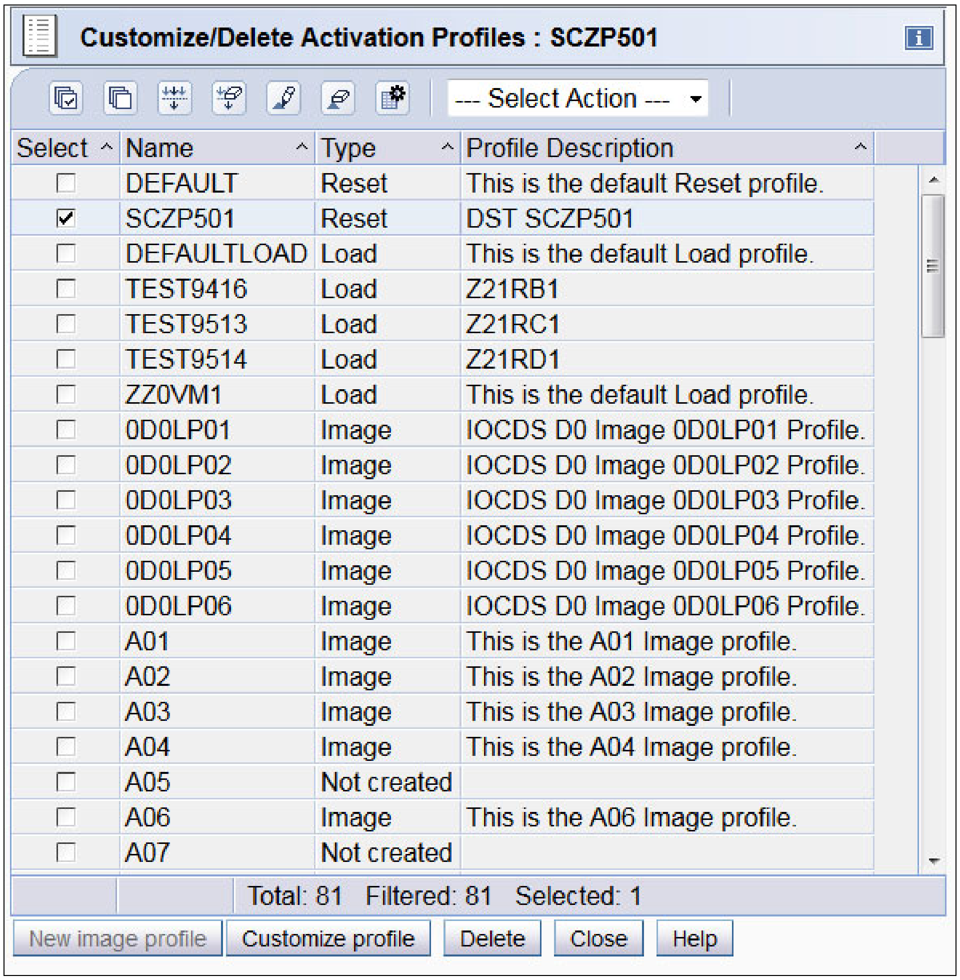

Select	Name	Type	Profile Description
☐	DEFAULT	Reset	This is the default Reset profile.
☑	SCZP501	Reset	DST SCZP501
☐	DEFAULTLOAD	Load	This is the default Load profile.
☐	TEST9416	Load	Z21RB1
☐	TEST9513	Load	Z21RC1
☐	TEST9514	Load	Z21RD1
☐	ZZ0VM1	Load	This is the default Load profile.
☐	0D0LP01	Image	IOCDS D0 Image 0D0LP01 Profile.
☐	0D0LP02	Image	IOCDS D0 Image 0D0LP02 Profile.
☐	0D0LP03	Image	IOCDS D0 Image 0D0LP03 Profile.
☐	0D0LP04	Image	IOCDS D0 Image 0D0LP04 Profile.
☐	0D0LP05	Image	IOCDS D0 Image 0D0LP05 Profile.
☐	0D0LP06	Image	IOCDS D0 Image 0D0LP06 Profile.
☐	A01	Image	This is the A01 Image profile.
☐	A02	Image	This is the A02 Image profile.
☐	A03	Image	This is the A03 Image profile.
☐	A04	Image	This is the A04 Image profile.
☐	A05	Not created	
☐	A06	Image	This is the A06 Image profile.
☐	A07	Not created	

Figure 11-12 The HMC: Customize/Delete Activation Profiles

3. In the reset profile, select the new IOCDS that was written to the server in the General page.
4. A window opens, asking if you want to maintain the logical partition activation order in the Partitions page. Remember to verify the activation order of the coupling facility partitions on the server if you do update the activation order. Failure to do so can cause coupling facility partitions to activate after partitions running in other modes. This might lead to operating systems entering a disabled wait state.

Defining the image profile to the new logical partition

You are now ready to define all the necessary resources to the new logical partition. After you activate the production IODF that contained the newly named logical partition, you can view that logical partition from the HMC. You now must build the image profile for that partition. For more information about the definitions in the image profile, see *zEnterprise System Processor Resource/Systems Manager Planning Guide*, SB10-7156.

To create the image profile for a new logical partition, select **Operational Customization** → **Customize/Delete Activation Profiles** in the HMC. Select an existing image profile and then click **Customize selected profile**, or click the **Image profile wizard**. Define the values for the new logical partition.

Details about the definition values, organized by page, are in the following sections.

General page

The following definition values are for the General page (Figure 11-13 on page 584):

- Partition identifier

  You can assign a partition identifier to the new logical partition. This is defined by a hex value that represents a maximum 85 (decimal) logical partitions. The valid identifiers for LPARs are X'0' through X'7F'. The LPAR identifier must be unique for each active LPAR. Remember to use a value that is not in use by any other partition. If the new logical partition is to be used as a coupling facility, the partition ID is important when defining or updating the CFRM policy. Generally, the partition ID should be a concatenation of CSSID and MIFID. For this example, the LPAR is MIFID 4 in CSS 1, so the partition ID is X'14'.

- Mode (of operation)

  You can define how this partition is to be used. The mode of an LPAR, depending on the model, can be ESA/390, ESA/390 TPF, Linux-Only, z/VM, IBM zAware, or coupling facility. The mode of an LPAR must support the mode of the control program that is loaded into it. ESA/390 LPARs support ESA/390 control programs. Coupling facility LPARs support the coupling facility control code. z/VM LPARs support z/VM, IBM zAware supports IBM zAware LICC, and Linux-Only LPARs support Linux or z/VM.

  > **zIIP information:** z Systems Integrated Information Processors (zIIPs) can be defined to ESA/390 or z/VM mode image. However, zIIPs are supported only by z/OS. Other operating systems cannot use zIIPs, even if they are defined to the logical partition. z/VM V5.3 or later can provide zIIPs to a guest z/OS.

- Clock Type Assignment

  With the clock value, you can decide to select the **Standard time of day**, or to set an offset value to the STP clock here if you want by selecting **Logical partition time offset**.

  When this option is selected, another option (Time Offset) is added to the logical partition activation profile. You can set the logical partition time either ahead of or behind your STP time by selecting the appropriate option.

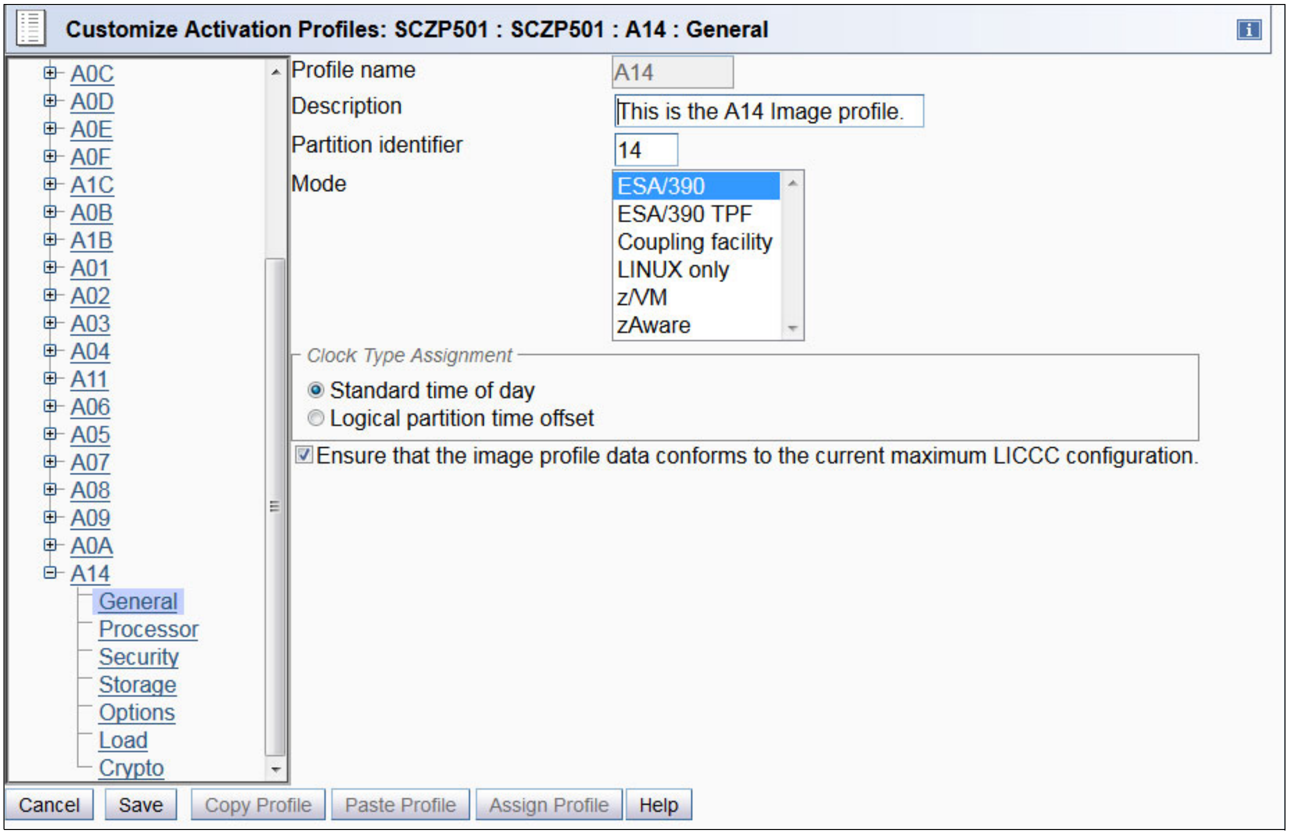

Figure 11-13 The HMC: General page in Image Profile

Processor page

The Processor page (Figure 11-14 on page 585) in the Image profile is where you specify the new logical partition CP assignment. The items that are selectable on this page change depending on the mode of the LPAR selected on the General page:

- Logical Processor Assignments

  Figure 11-14 on page 585 shows a sample page for an ESA/390 mode logical partition. You must set the initial and reserved number of CPs and zIIPs. zAAPs are not available any more on z13. *Initial* refers to how many processors come online when an IPL is performed on a logical partition. *Reserved* refers to how many processors can be configured online while the system is running. This value also allows for future upgrades, and can be set to more than what is physically available in the server. That way, when extra processors are enabled, they can be configured online and used immediately.

- Not Dedicated Processor Details

  If you operate this logical partition with *shared* processors (the **Dedicated processors** check box in the Logical Processor Assignments section is *not* selected), you must set the processing weight and, if needed, the initial capping value and the Workload Manager values. Consider the impact on the systems that are working in other logical partitions on the same processor. The effective capacity for a processor might decrease because of the cost to manage logical partitions becomes larger for PR/SM. Estimate the capacity by using the zPCR tool before you define processor resources, including the weight value. You can obtain this tool from IBM. For more information, see the IBM Techdoc *Getting Started with zPCR (IBM Processor Capacity Reference)*, PRS1381.

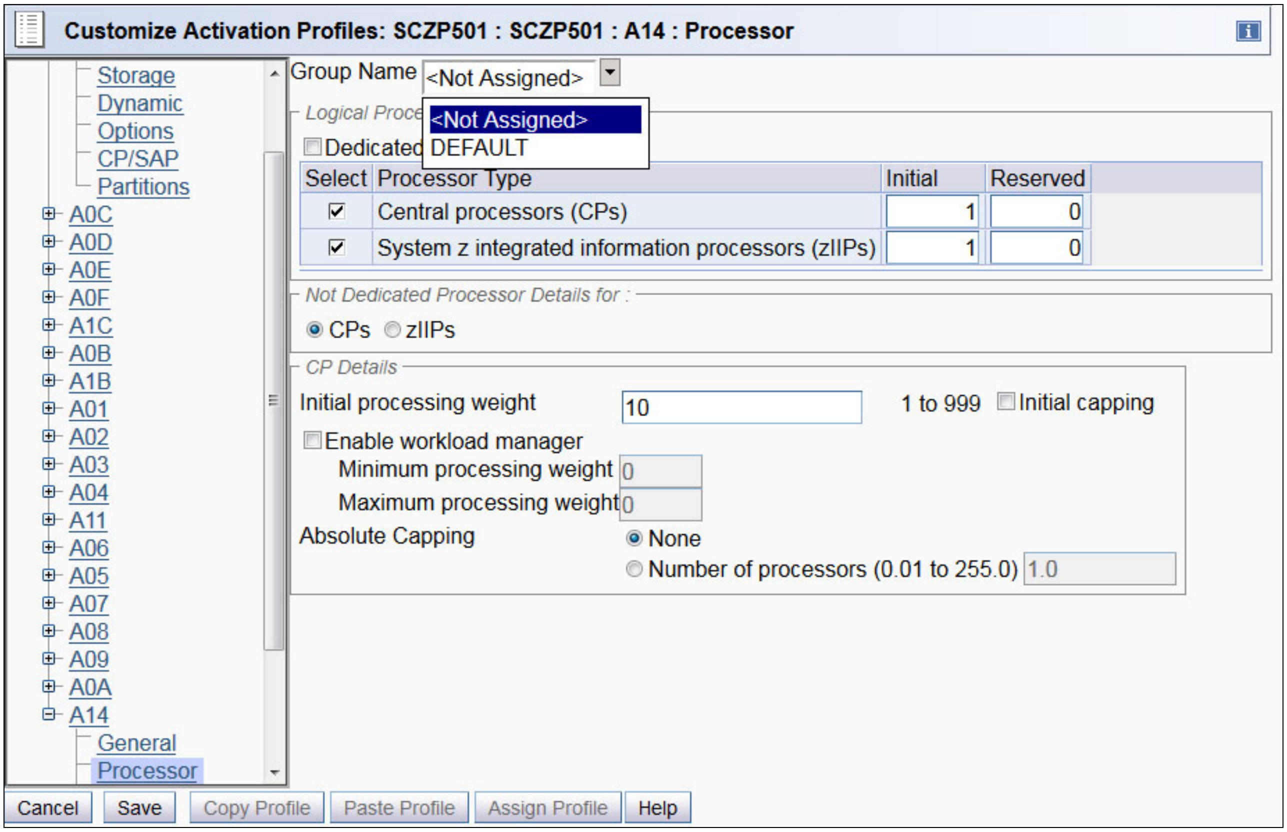

Figure 11-14 The HMC Processor page in image profile

Security page

The Security settings determine the level of interlogical partition action that occur. These values are unique to each client, but generally the first two options are always selected:

- Global performance data control
- Input/Output configuration control

In addition, these options can be selected:

- Cross partition authority: To indicate whether the logical partition can be used to issue control program instructions that reset or deactivate other logical partitions.
- Logical partition isolation: Use for greater partition isolation,

Storage page

Depending on the hardware model, the z13 can support a minimum of 64 GB to a maximum of 10 TB of memory for client use, and up to 1 TB per LPAR for z/OS V1.12 and V1.13, and up to 4 TB per LPAR for z/OS V2.1. Before you activate the new logical partition, confirm the size of available storage to assign to the new logical partition. From the SE, select the **Operational Storage** → **Storage Information** task to determine the available memory (Figure 11-15).

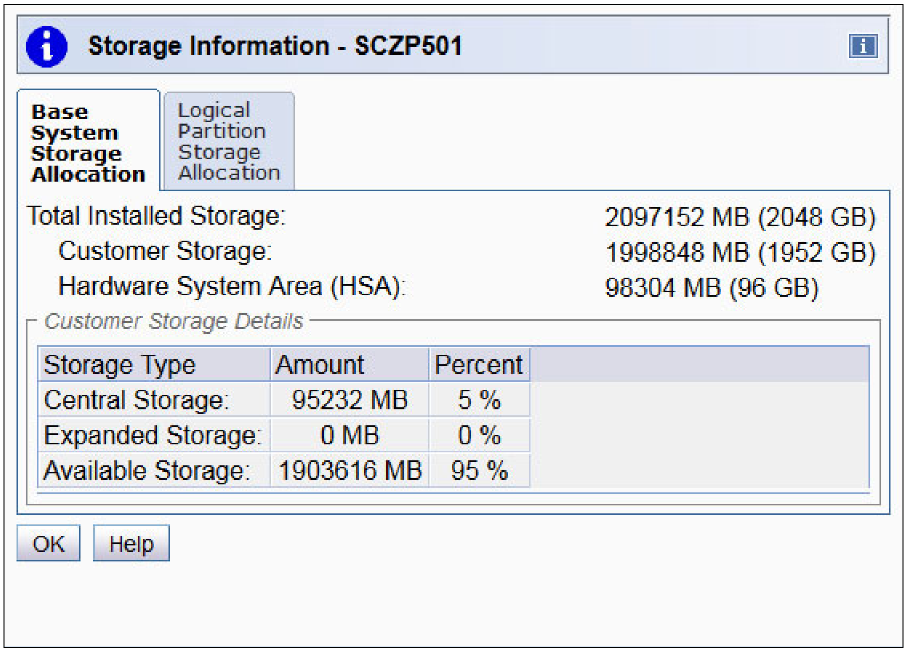

Figure 11-15 SE: Storage Information

This task also shows that the HSA size is a fixed value (96 GB). The sum of the main and expanded storage that you define in the image profile must be set to a value smaller than the amount of available storage, as shown in the Storage Information window.

Storage options are set from the Storage page in the image profile (Figure 11-16 on page 587). On the z13, memory can be assigned as a combination of main storage and expanded storage, supporting up to 85 logical partitions. Enter the value (in megabytes) of each storage type.

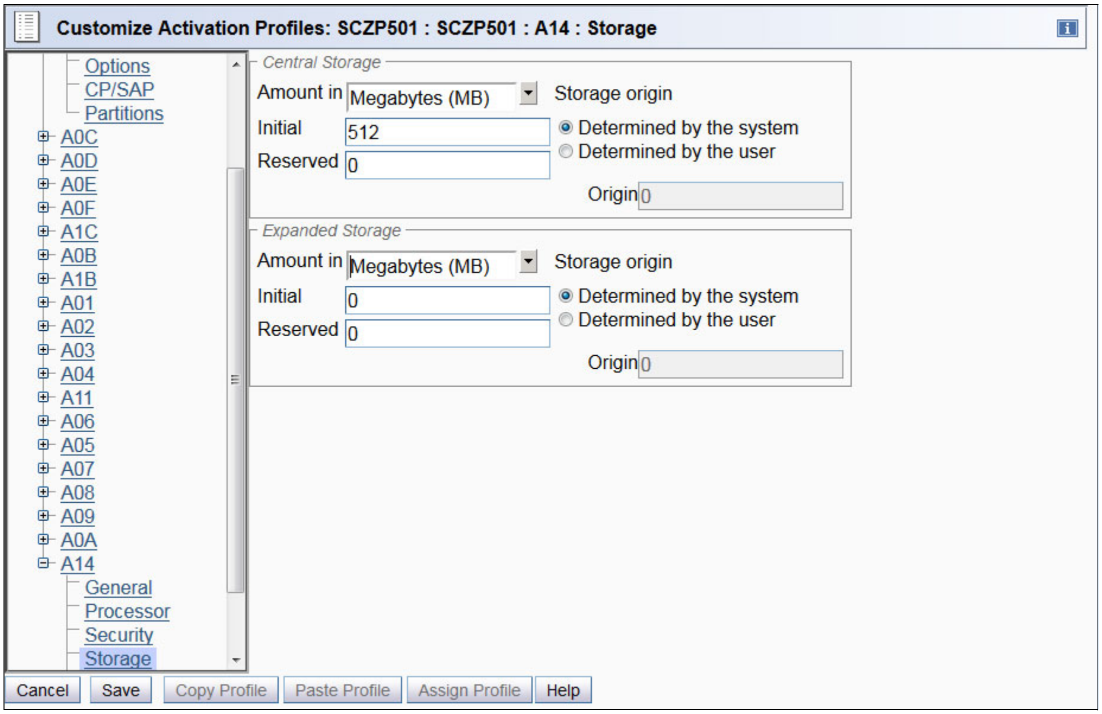

Figure 11-16 The HMC: Storage page in image profile

Initial storage is the amount of storage to be allocated to the LPAR at activation.

Reserved storage is the amount of extra storage that is requested for the LPAR. Reserved storage is storage that can be dynamically brought online to the LPAR at some point after LPAR activation. Entering 0 for reserved storage limits storage to the initial amount during the LPAR activation.

Granularity of main storage for a logical partition depends on the largest main storage amount that is defined for either initial or reserved main storage, as shown in Table 11-1. The granularity applies across all main storage that is defined, both initial and reserved. For example, for an LPAR with an initial storage amount of 60 GB and a reserved storage amount of 160 GB, the main storage granularity of both initial and reserved main storage is 1 GB. Expanded storage granularity is fixed at 512 MB. Expanded storage granularity applies across the expanded storage input fields.

Table 11-1 Central Storage granularity

LPAR main storage amount	LPAR storage granularity
512 MB <= main storage amount <= 256 GB	512 MB
256 GB < main storage amount <= 512 GB	1 GB
512 GB < main storage amount <= 1 TB	2 GB
1 TB < main storage amount <= 2 TB	4 GB
2 TB < main storage amount <= 4 TB	8 GB
4 TB < main storage amount <= 8 TB	16 GB
8 TB < main storage amount <= 10 TB	32 GB

Both an IBM z/Architecture (64-bit) mode and an ESA/390 (31-bit) architecture mode operating system can run in an ESA/390 image on a z13. Any ESA/390 image can be defined with more than 2 GB of main storage and can have expanded storage. These options allow you to configure more storage resources than the operating system can address.

Cryptographic options

If the newly named logical partition is using cryptographic services, you must enable these options in the Crypto page of the image profile. Make sure that the control domain index and usage domain index are correct, and do not conflict with any other logical partitions on the z13. If you duplicate definition values across logical partitions, you cannot save the logical partition activation profiles until you correct the cryptographic definitions.

These values can be defined by using the Change LPAR Cryptographic Controls task from the SE after you activate the logical partition using the new image profile. For more information, see 10.2.3, "Cryptographic configuration using the Support Element (SE)" on page 550.

Activating a defined logical partition

After all the values are updated and saved, the logical partition can be activated and an IPL can be run. This operation is done from the HMC by completing the following steps:

1. Select the logical partition from the HMC Contents of the CPC that you want to activate. Click **Activate** from the Daily group menu.
2. In the Activate Task Confirmation window (Figure 11-17), confirm that the activation profile is correct, and click **Yes**.

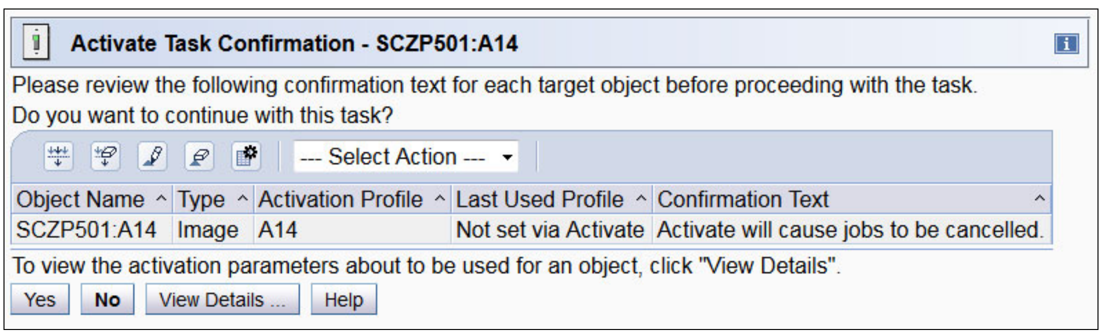

Figure 11-17 The HMC: Activate Task confirmation

3. Verify that the activation task completed successfully. The status changes from `Not activated` to `Not operating` in the detail information of the target logical partition (Figure 11-18).

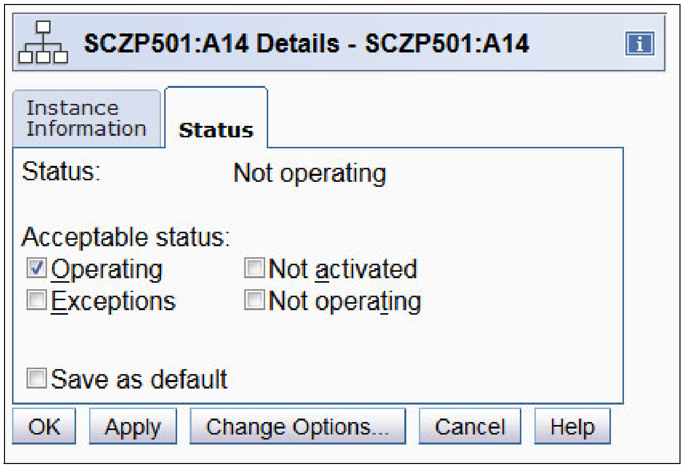

Figure 11-18 The HMC: Details after you activate a new logical partition

Finally, you can run an IPL of the operating system in the newly named logical partition.

> **Remember:** Any HMC users that are restricted to specific LPARs by Customized User Controls Managed Resources cannot access the newly added LPAR.

11.3 Changing a logical partition to reserved

This section explains the steps to change an active logical partition into a reserved partition whose name is removed (unnamed). It does not describe how to *delete* a logical partition. You cannot delete a logical partition (LP).All 85 LPs are either defined (named) or reserved (unnamed).

Deactivating a logical partition

First, you must deactivate the target logical partition. Deactivating a logical partition is a straightforward process. After the operating system is reset, the logical partition can be deactivated by using the Deactivate task in the HMC. This action frees all the partition's resources, such as memory, CPs, and channels.

HCD definition changes

After deactivation is complete, make changes (by HCD or HCM) to revert the logical partition to a reserved logical partition. In HCD, the process of changing a named logical partition to a reserved logical partition is the reverse of naming a logical partition. Before you make any changes in HCD, verify which channels and devices are defined to the Candidate Lists for this logical partition.

Some devices might generate errors if they have been included in the candidate lists for the logical partition that is changed to reserved status.

To change the HCD definition, complete the following steps:

1. In the HCD panel, select the CSS ID that contains the partition that you want to change to a reserved logical partition. A sample partition list is shown in Figure 11-1 on page 575.

   Important: First, remove all assigned CHPIDs for this logical partition. Failure to do so before changing the logical partition name to an asterisk (*) results in HCD generating errors.

2. Select the **Change Partition Definition** option and proceed to the Update CHPID Access and Candidate Lists panel (Figure 11-3 on page 576). Change all the CHPID Access and Candidate Lists to `No`. This process removes all channels from the logical partition.
3. In the Change Partition Definition panel (Figure 11-19), rename the logical partition to an asterisk (*), set the Partition usage field back to `CF/OS`, and delete any data in the Description field. You are now ready to activate the new configuration.

```
------------- Change Partition Definition --------------


Specify or revise the following values.

Partition name . . . *
Partition number . . 4      (same as MIF image ID)
Partition usage  . . CF/OS  +

Description  . . . . ____________________
```

Figure 11-19 Changing a partition to reserved status

Dynamic activation of configuration

To dynamically activate the production IODF, complete the following steps:

1. Select option **1. Activate new hardware and software configuration**, as shown in Figure 11-9 on page 580. When a new configuration that has a reserved logical partition (unnamed) is activated, you must add the force options. If the force options are not specified, the activation fails with an IOS500I message (Figure 11-20).

```
------------------------------- Message List --------------------------------
   Save  Query  Help
 --------------------------------------------------------------------------
                                                                 Row 1 of 15
 Command ===> ___________________________________________ Scroll ===> PAGE

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

 / Sev Msg. ID  Message Text
 _ I   IOS500I  ACTIVATE RESULTS
 #              TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN
 #              ACTIVATE FAILURE
 _ I            NOTE = 0112,REQUEST CONTAINS DELETE(S), SPECIFY FORCE ON
 #              H/W ACTIVATE
 #               COMPID=SC1C3
 _ I   CBDA888I Following devices are to be modified for processor
 #              SCZP501: 1.2100-1.2108,1.210F,1.2120-1.2128,1.212F
 _ I   CBDA882I Following channel paths are to be modified for processor
 #              SCZP501:
 #              1.00-1.02,1.07-1.09,1.0B-1.0D,1.10-1.13,1.40-1.48,1.4F-1.
 #              53,1.60-1.68,1.F3-1.F4,1.FE-1.FF
_ I   CBDA878I Following partitions are to be deleted from processor
 #              SCZP501: 1.A14
 _ I   CBDA126I ACTIVATE IODF=42,TEST command was accepted.
```

Figure 11-20 HCD: Activate new IODF remove logical partition with message

2. Set the **Allow hardware deletes** and the **Delete partition access to CHPIDs unconditionally** options to `Yes` (Figure 11-21).

```
------------- Activate New Hardware and Software Configuration -------------


 Specify or revise the values for IODF activation.

 Currently active IODF . : SYS6.IODFE7
    Processor ID . . . . : SCZP501      Sphinx
    Configuration ID . . : L06RMVS1     Sysplex systems
    EDT ID . . . . . . . : 01

 IODF to be activated  . : SYS6.IODF42
    Processor ID . . . . . SCZP501   +
    Configuration ID . . . L06RMVS1  +     EDT ID . . . 01  +

 Test only  . . . . . . . . . . . . . . . . . . . . . No   (Yes or No)
 Allow hardware deletes (FORCE, FORCE=DEVICE) . . . . Yes  (Yes or No)
 Delete partition access to CHPIDs unconditionally
 (FORCE=CANDIDATE)  . . . . . . . . . . . . . . . . . Yes  (Yes or No)
 Write IOCDS  . . . . . . . . . . . . . . . . . . . . No   (Yes or No)
 Switch IOCDS for next POR  . . . . . . . . . . . . . No   (Yes or No)
```

Figure 11-21 Activate New Hardware and Software Configuration

3. Review all messages in the resulting panel (Figure 11-22) and verify that the dynamic activation completed successfully.

```
------------------------------ Message List --------------------------------
   Save  Query  Help
 --------------------------------------------------------------------------
                                                                Row 1 of 12
 Command ===> ___________________________________________ Scroll ===> PAGE

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

 / Sev Msg. ID  Message Text
 _ I   IOS500I  ACTIVATE RESULTS
 #              ACTIVATE COMPLETED SUCCESSFULLY
 _ I   CBDA888I Following devices are to be modified for processor
 #              SCZP501: 1.2100-1.2108,1.210F,1.2120-1.2128,1.212F
 _ I   CBDA882I Following channel paths are to be modified for processor
 #              SCZP501:
 #              1.00-1.02,1.07-1.09,1.0B-1.0D,1.10-1.13,1.40-1.48,1.4F-1.
 #              53,1.60-1.68,1.F3-1.F4,1.FE-1.FF
 _ I   CBDA878I Following partitions are to be deleted from processor
 #              SCZP501: 1.A14
 _ I   CBDA126I ACTIVATE IODF=42,FORCE=(DEVICE,CANDIDATE) command was
 #              accepted.
***************************** Bottom of data ******************************
```

Figure 11-22 HCD: Activate new IODF remove logical partition with message

The activation action run by HCD can also be accomplished by issuing the **`activate iodf=xx,force=(candidate,device)`** z/OS command.

Removing resources from the activation profiles

When the new configuration has been activated and the logical partition changed to reserved, you can remove any profiles that are defined for the logical partition from the HMC.

After all the resources are removed, write a new IOCDS to the SE on the server and make it active. Afterward, update the Activation profiles to point to the new IOCDS. The partition that was changed to reserved no longer exists in the partition list.

Consideration: Remember to save the Reset activation profile in case a recovery action must be run against the server. If the old IOCDS is still in the Reset profile and a recovery action is run against the server, the old partition is defined to the server again and attempts to allocate all its resources.

11.4 Renaming a logical partition

Changing a logical partition name from one name to another name cannot be done directly. It is a two-step approach that combines the following tasks:

1. Deactivating and unnaming a named logical partition, making it reserved.
2. Naming, assigning resources, and activating the reserved logical partition.

Consider the following information when you use that approach:

1. The configuration must be changed to remove all assigned channels and devices, the partition name must be set to asterisk (*), and the new configuration must be activated.
2. The configuration must be changed again and the reserved logical partition must be given a new name (and, if need be, a new operating mode). All the channels and devices that it requires must be assigned.

   The configuration is then activated and the newly named logical partition is ready for use.

For more information, see the following sections:

- To unname and deactivate a logical partition, see 11.3, “Changing a logical partition to reserved” on page 589.
- To name a reserved logical partition, see 11.2.1, “Naming a reserved logical partition” on page 575.

This section concentrates on changing the name and operating mode of a logical partition. The example used is changing a z/OS partition to a Linux partition.

Any of the following combinations are possible:

- z/OS partition to CF partition, or to a Linux partition, or to a different z/OS partition
- CF partition to z/OS partition, or to a Linux partition, or to a different CF partition
- Linux partition to a z/OS partition, or to a CF partition, or to a different Linux partition

Renaming a z/OS logical partition to a Linux logical partition

Changing a named z/OS partition to a Linux partition consists of six steps. The assumption is that all the resources needed for the Linux partition are available (including storage, channels, and IFLs or CPs).

To accomplish this change, complete the following steps:

1. In the IODF, remove all assigned channels and, optionally, any devices that might be in the candidate list for the logical partition that will have its name and usage changed.
2. Change the partition name value to an asterisk (*). This action makes the logical partition a reserved logical partition. Build the new IODF. Make sure the logical partition that is renamed has been reset and deactivated on the HMC application.
3. Dynamically activate the new configuration. You can use either the **`activate`** command from an active z/OS logical partition on the z13 where the logical partition is being renamed, or the HCD.
4. Using the work version of the IODF activated in step 3, name the reserved logical partition. Assign a partition name and usage. In this example, the usage can be CF/OS, but OS is also allowed. Keep in mind that the partition MIFID *cannot* be changed.

   > **Tip:** If you are renaming a partition on a z13 that already has a coupling facility with coupling links defined, change the partition usage value to OS. Coupling links can be connected to only one coupling facility on a server.
   >
   > By using CF/OS, HCD prevents you from connecting any coupling link that is already connected to another coupling facility to the reserved partition you are renaming. This occurs even if you are using this logical partition as an operating system.

5. Assign the channel paths and devices that are needed for the Linux logical partition. These can include channel types CIB, CS5, FC, FCP, ICP, IQD, OSC, OSD, OSE, OSM, OSN, and OSX.
6. You can now build a new production IODF for dynamic activation. The dynamic activation can be done again by using the **`activate`** z/OS command, or by using HCD.

   After the activation is successful, write a new IOCDS to the SE. The IOCDS in the Activation Profile for the z13 still reflects the old name of the logical partition. Although this activation profile for the renamed logical partition can be updated and used, the name might not match the new name you assigned. Change the activation profile to point to the new IOCDS and pay attention to the activation order of all the logical partitions.

12

Preparing IBM z13 for zBX Model 004

The zBX is designed to integrate multi-platform system and heterogeneous workloads that are built with IBM certified components. The advanced virtualization management capabilities of the IBM z Systems (offered through the Unified Resource Manager) allow the zBX to be managed as a single pool of resources.

The zBX (Machine type 2458) Model 004 is an independent member of an ensemble, no longer controlled by an owning CPC.

For more information about the role of zBX in the z Systems environment and the Unified Resource Manager, see the material listed in 12.7, "Additional references" on page 629.

This chapter introduces the zBX Model 004 and information to consider when you plan to convert an existing zBX Model 002 or Model 003 to a Model 004.

This chapter includes the following topics:

- zBX hardware summary
- zBX connectivity in an ensemble
- Upgrade considerations
- Upgrade overview
- Configuration examples
- Ensembles and Unified Resource Manager
- Additional references

Information: For zBX installation and setup, see 14.3, "HCD definitions and z/OS display commands for zBX" on page 672.

12.1 zBX hardware summary

The zBX Model 004 (Upgrade Feature Code 0512) is an independent node in an ensemble. It has its own two Support Elements (SE) and unlike prior zBX Models, it is no longer controlled by a CPC. Having additional nodes in the ensemble to operate the zBX Model 004 is unnecessary. A zBX Model 004 cannot be ordered as a new system. Only an existing zBX Model 002 or Model 003 can be upgraded to a zBX Model 004.

If a controlling z196 or zEC12 is upgraded to an IBM z13, the attached zBX must be upgraded to a zBX Model 004.

All blades of Model 002 and Model 003, with the exceptions of ISAOPT blades, are carried forward and supported on the Model 004.

> **Disruptive:** Upgrades from zBx Model 002 or Model 003 to Model 004 are disruptive.

The new zBX Model 004 can consist of the following key components:

- Up to four standard 19-inch 42U IBM Enterprise racks with required power infrastructure
- One to eight BladeCenter chassis (one or two per rack) with a combination of up to 112[1] different blades. The maximum number of blades per feature in a zBX with eight BladeCenter chassis can be up to these amounts:
  - 28 IBM DataPower® blades (Feature Code 0611). DP blades need two slots inside a BladeCenter Chassis. This means seven DataPower Blades can be plugged in a single BladeCenter Chassis (fourteen DP blades per frame)
  - 56 xASB Blades (Feature Code 0613).
  - 112 pASB Blades (Feature Code 0612).
- Two Advanced Management Modules (AMM) in each BladeCenter chassis to manage and monitor all components in the BladeCenter chassis
- Two 8-Gbps Fibre Channel switch modules in each BladeCenter chassis to connect to customer's storage area network (SAN)
- Two 1U SEs in the first rack (Rack B)
- Four VLAN capable Top of Rack switches in the first rack (Rack B):
  - Two switches for CPC to zBX (OSX) connections, zBX to zBX connections, and zBX connections to customer's network
  - Two intranode management network (INMN) switches connecting to the SEs of the zBX
- Two keyboard-video-mouse (KVM) trays in Rack B to access the two SEs
- Optional acoustic rear door or optional rear door heat exchanger

The new zBX Model 004 includes the following improvements:

- No need for a direct INMN link between CPC and zBX. This eliminates the restriction of the maximum distance of 26 meters to a owning CPC.

[1] The number of chassis and blades varies depending on types of the blades that are configured within zBX.

- The zBX Model 004 is now a stand-alone node in the ensemble, eliminating any interferences with CPC operations and vice versa. This applies to these items:
  - Hardware maintenance.
  - System upgrade.
  - Code load.
  - Power On/Off scenarios.
  - There is no need to access a CPC's SE (physically, or remotely for any zBX related task).
  - There is no load on the CPC's SE caused by zBX Operations and vice versa.
  - The management of CPC resources and zBX resources is always routed through their dedicated SEs (no access to the zBX SE to manage a CPC resource and vice versa).

Power consumption: A zBX Model 004 Rack B always requires four connections to wall power. A zBX Model 003 with a single BladeCenter chassis in Rack B required only two connections to wall power. There is a now higher power consumption because of the extra SEs and KVM trays.

Figure 12-1 shows a zEC12 Ensemble with a zEC12 and zBX Model 003 attached, and a z196 with an attached zBX Model 002.

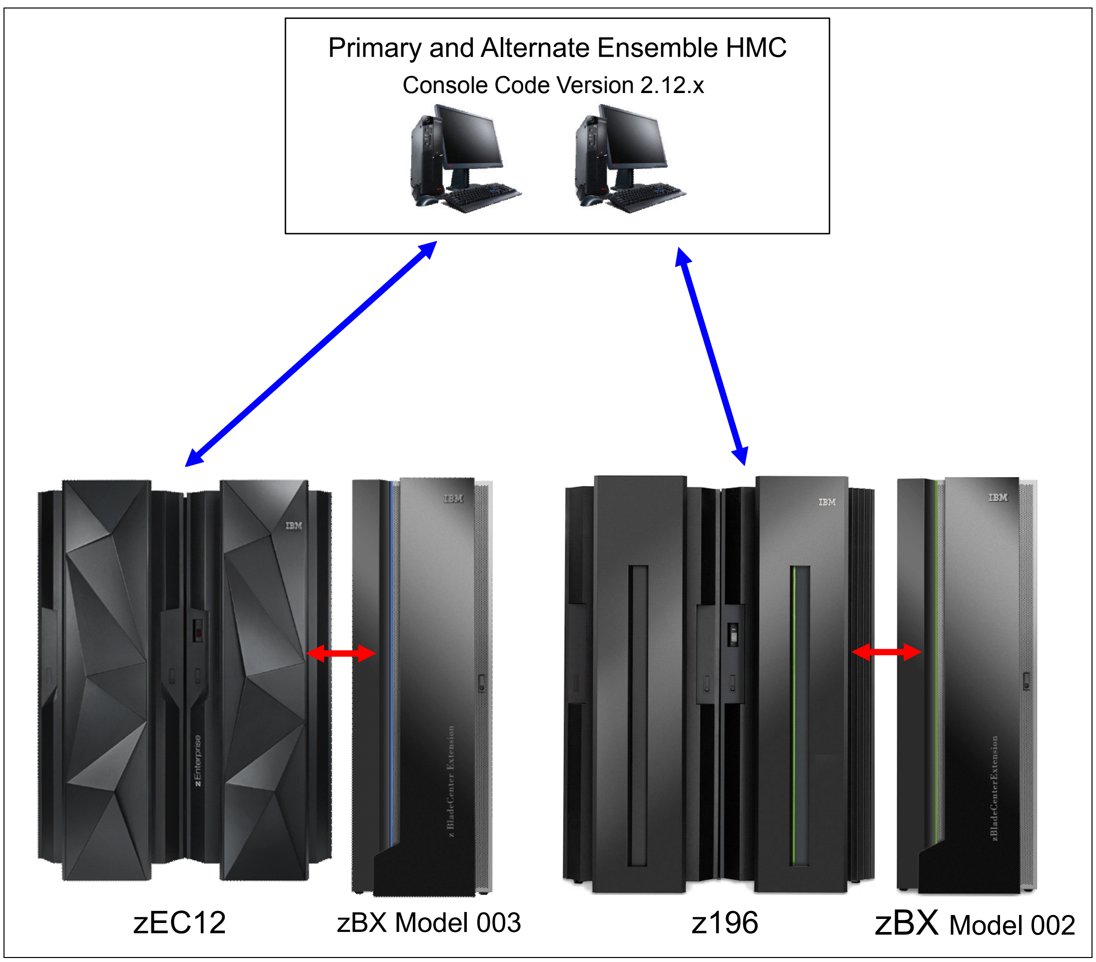

Figure 12-1 zEC12 Ensemble

Figure 12-2 shows an IBM z13 Ensemble with four independent nodes: one z13, one z196, and two zBX Model 004 upgraded stand-alone models.

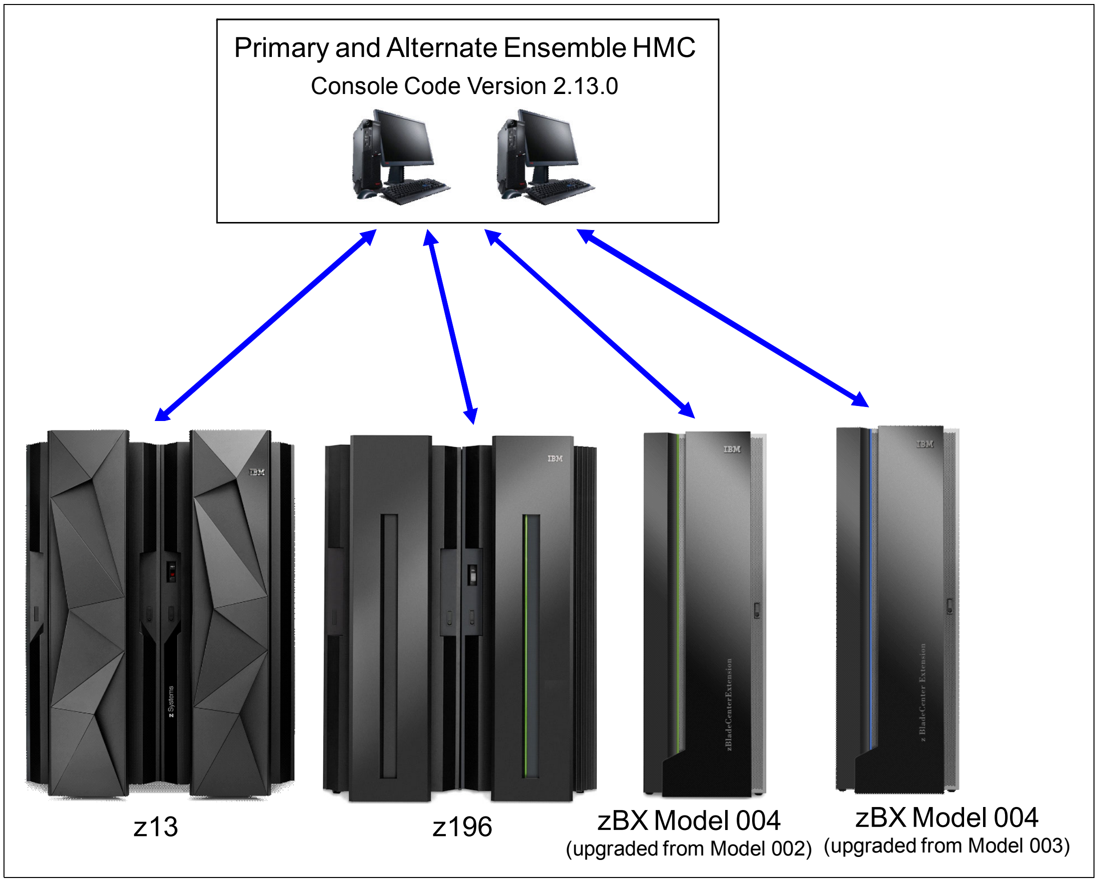

Figure 12-2 IBM z13 Ensemble

The Rack B is the primary rack. It contains up to two BladeCenter chassis that can contain up to 14 blades each, two SEs, two KVM Modules (a 1U tray containing a keyboard and a display dedicated to one SE), and four Top of Rack (TOR) switches to support the intranode management network (INMN) and the intraensemble data network (IEDN). These networks are used for the overall operation and support of the zEnterprise System. For more information, see 12.2.2, “Intranode management network (INMN)” on page 604, and 12.2.3, “Intraensemble data network (IEDN)” on page 609.

Racks C, D, and E are expansion racks. They can accommodate more BladeCenter chassis. Figure 12-3 on page 599 shows the order of BladeCenter chassis as they are ordered and installed. Also, notice from the front of the zBX that the TOR switches have an installed cover plate. All cabling is done from the rear of the system.

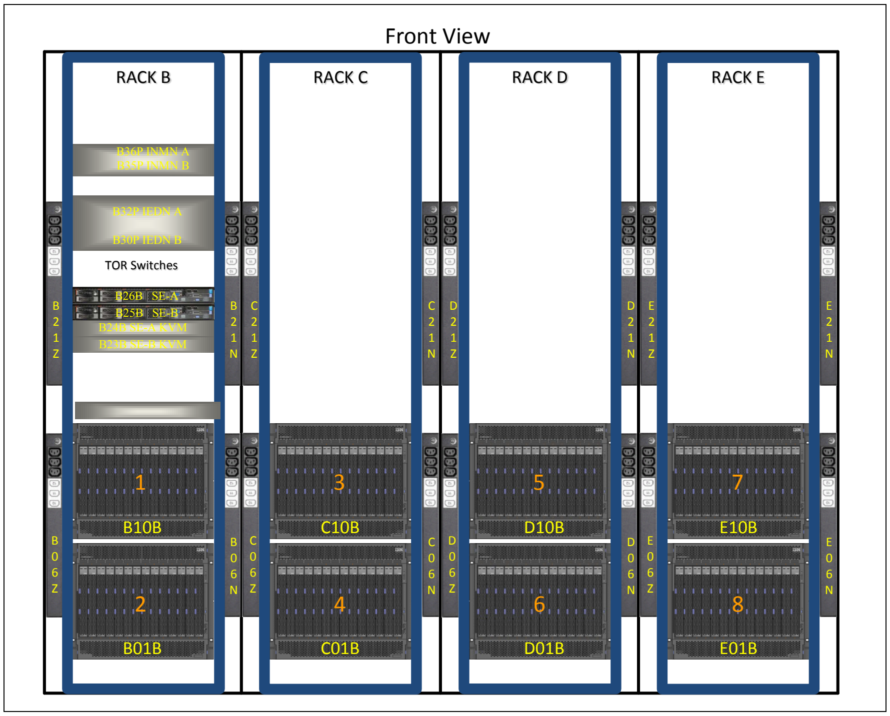

Figure 12-3 BladeCenter chassis: Order in which the chassis are installed in the racks

Figure 12-4 shows the rear view of the four TOR switches installed in Rack B. From top to bottom are two 1000BASE-T switches for the INMN and two 10 GbE switches for the IEDN. The INMN and IEDN TOR switches are cabled for redundancy. Power distribution units are installed in the rear of the racks to provide power for the internal components.

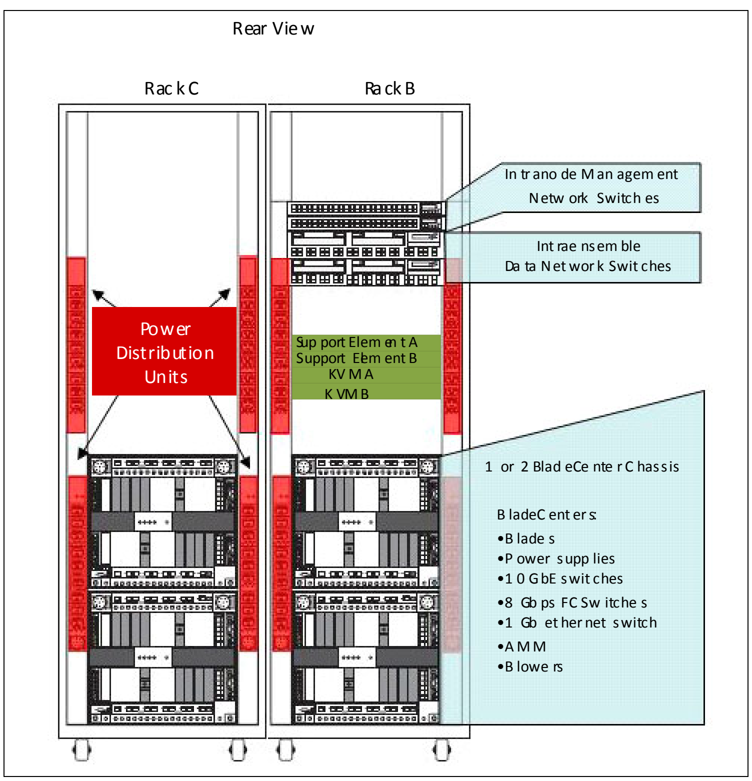

Figure 12-4 zBX racks: Rear view

Racks C through D require 2 - 4 independent redundant client-supplied power feeds. Rack B always requires four independent power feeds.

For more information about power and environmental requirements, see these publications:

- *Installation Manual for Physical Planning*, GC27-2630
- *Installation Manual - Physical Planning 2827,* GC28-6914

The zBX includes the following optional benefits.

- The IBM Rear Door Heat eXchanger (FC 0540) reduces the heat load of the zBX emitted into ambient air. The Rear Door Heat eXchanger is an air-to-water heat exchanger that diverts heat of the zBX to chilled water (client supplied data center infrastructure). The Rear Door Heat eXchanger requires external conditioning units for its use.
- Height reduction (FC 0570) reduces the rack height to 36U high and accommodates doorway openings as low as 1832 mm (72.1 inches). Order this choice if the openings of your doorways are less than 1941 mm (76.4 inches) high.

When the zBX is upgraded to Model 004, you may order the following features for a defined period of time:

- An entitlement record allowing the installation of more IBM System x blades and IBM POWER7® PS701 blades, up to the limit of available empty (not used) slots in the existing BladeCenter chassis
- Additional Small Form Factor Pluggable (SFP) optics to change or extend the zBX connectivity
- Top Exit Support (FC 0601 for Rack B and FC 602 for Rack C through E)
- Hardware Management Consoles (HMC) FC 0092 and 0094 (FC 0091 can be carried forward)
- A change of DataPower Software features
- The Automate Firmware Suite (FC 0020). All nodes in the IBM z13 ensemble must be at the same level.
- The IBM Rear Door Heat eXchanger (FC 0540)
- The IBM Acoustic Door (FC 0543) to reduce the noise from the zBX

Note: DataPower blade moves are not allowed during model upgrade MES to a model 004. See 12.3, “Upgrade considerations” on page 616 for more information.

IEDN related length limitations by the fiber optic cable specifications and the connecting features valid for zBX Model 002 and Model 003 also apply for the zXB Model 004. It uses either short reach (SR) or long reach (LR) transceivers to terminate the IEDN. The fiber optic type must be specified when you order the zBX.

To establish connectivity either OSA-Express4S 10 GbE (carry forward only) or OSA-Express5S 10 GbE features can be used in the IBM z13 using CHPID type OSX. All fiber connections between the CPC and zBX for the IEDN are point-to-point.

Cable length restrictions for the intraensemble data network (IEDN) - 10 Gb Ethernet are shown in Table 12-1 on page 601.

Table 12-1 IEDN: Maximum lengths

Fiber optic cable types	Maximum distance
Single mode 9 micron for LR transceivers	10 km (6.2 miles)
Multimode 50 micron (2000 MHz-km) for SR transceivers	300 meters (984 feet)
Multimode 50 micron (500 MHz-km) for SR transceivers	82 meters (269 feet)
Multimode 62.5 Micron (200 MHz-km) for SR transceivers	33 meters (108 feet)

12.2 zBX connectivity in an ensemble

An ensemble is defined with the scope of platform management, and consists of a collection of one or more members.

An IBM z13 ensemble can include any of these nodes:

- IBM zEnterprise 196 (with or without zBX Model 002 attached)
- IBM zEnterprise 114 (with or without zBX Model 002 attached)
- IBM zEnterprise EC12 (with or without zBX Model 003 attached)
- IBM zEnterprise BC12 (with or without zBX Model 003 attached)
- IBM z13
- IBM BladeCenter Extension (zBX) Model 004

Figure 12-5 shows all nodes that can be present in an IBM z13 ensemble.

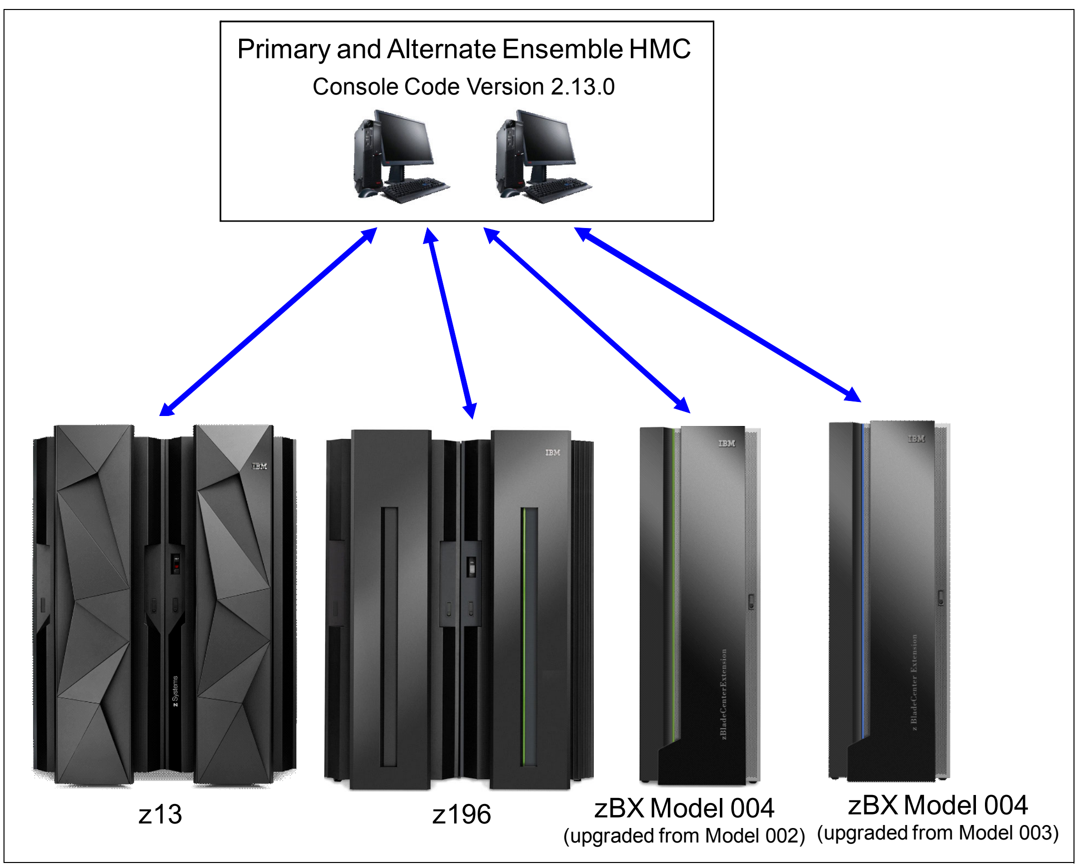

Figure 12-5 Overview of valid nodes in an IBM z13 ensemble

To participate in an ensemble, a CPC requires Feature Code 0025 and OSA-Express for Unified Resource Manager (OSM).

An ensemble must consist at least one of these nodes:

- A single CPC (two or more CPCs in an ensemble requires a zBX to be present)
- A single zBX Model 004
- A single CPC with a zBX Model 002/003 attached

An ensemble can have up to 16 nodes in its maximum configuration. This can be 1 - 8 CPCs and up to 8 zBXs.

A CPC in combination with the zBX consolidates a distributed business solution that is an interconnected set of applications. Some of these applications run on z Systems, and some run on System p or System x. A key value of an ensemble is that it brings all these applications together so they can be managed as a single unit through Unified Resource Manager.

The Unified Resource Manager provides advanced end-to-end management capabilities for both zBX and CPC. All components are configured, managed, and serviced the same way.

An ensemble consists of these networks:

- Customer provided management network
- Customer provided data network
- Intranode management network
- Intraensemble ensemble data network

Figure 12-6 shows a high level overview of all network connectivity to the ensemble and inside the ensemble.

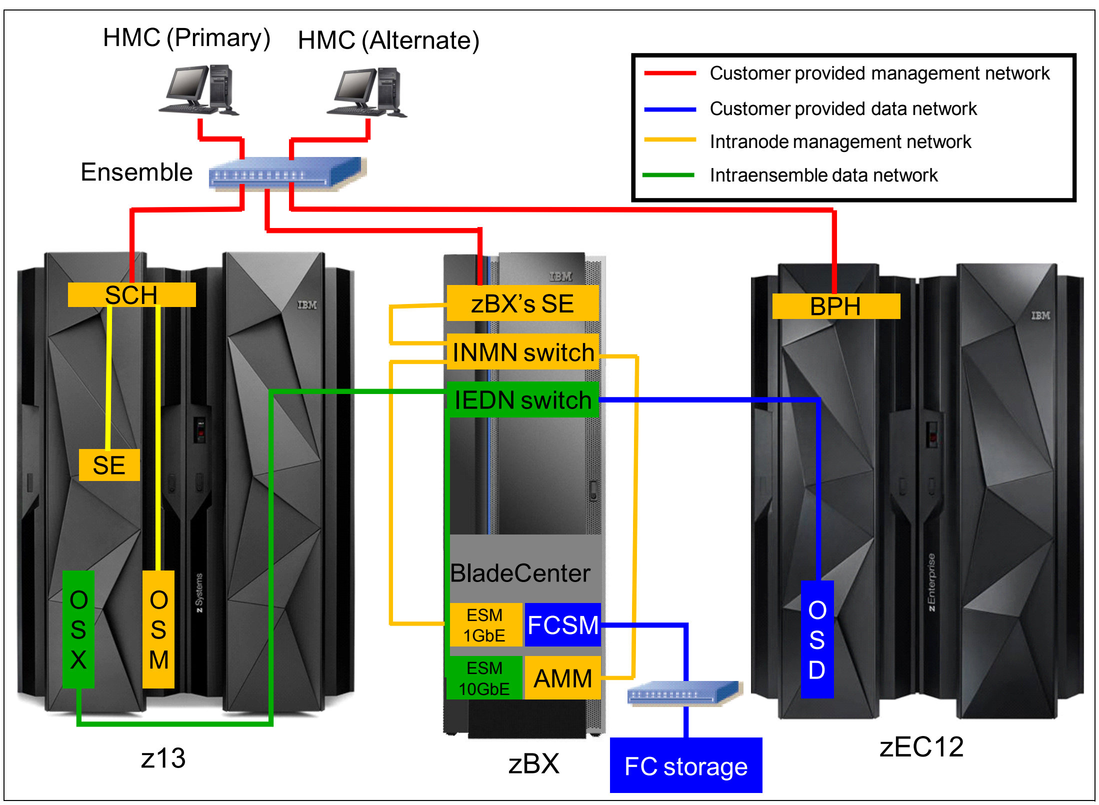

Figure 12-6 Overview showing all network connectivity related to the ensemble

12.2.1 HMC connectivity

The ensemble starts with a pair of HMCs that are designated as the primary and alternate HMCs and are assigned an ensemble identity. The CPCs and zBX systems are then added to the ensemble as members through an explicit action at the primary HMC.

The role of the alternate HMC is to mirror ensemble configuration and policy information from the primary HMC.

Both HMCs must be the same machine type/feature code and on the same LAN segment. This configuration allows the alternate HMC to take over the IP address of the primary HMC during failover processing.

Tip: The version of the console application for IBM z13 and zBX Model 004 is 2.13.0. The minimal feature code for both is 0091. HMC must have at least 16 GB of installed memory.

The communication path from an HMC to zBXs SE can consist of a complex network to resolve geographic or security issues. This HMC network is provided and managed by the client. For simplicity, only the necessary components and the communication path shown in Figure 12-7 are described:

- Dual HMCs configuration
- SE in zBX B-Frame

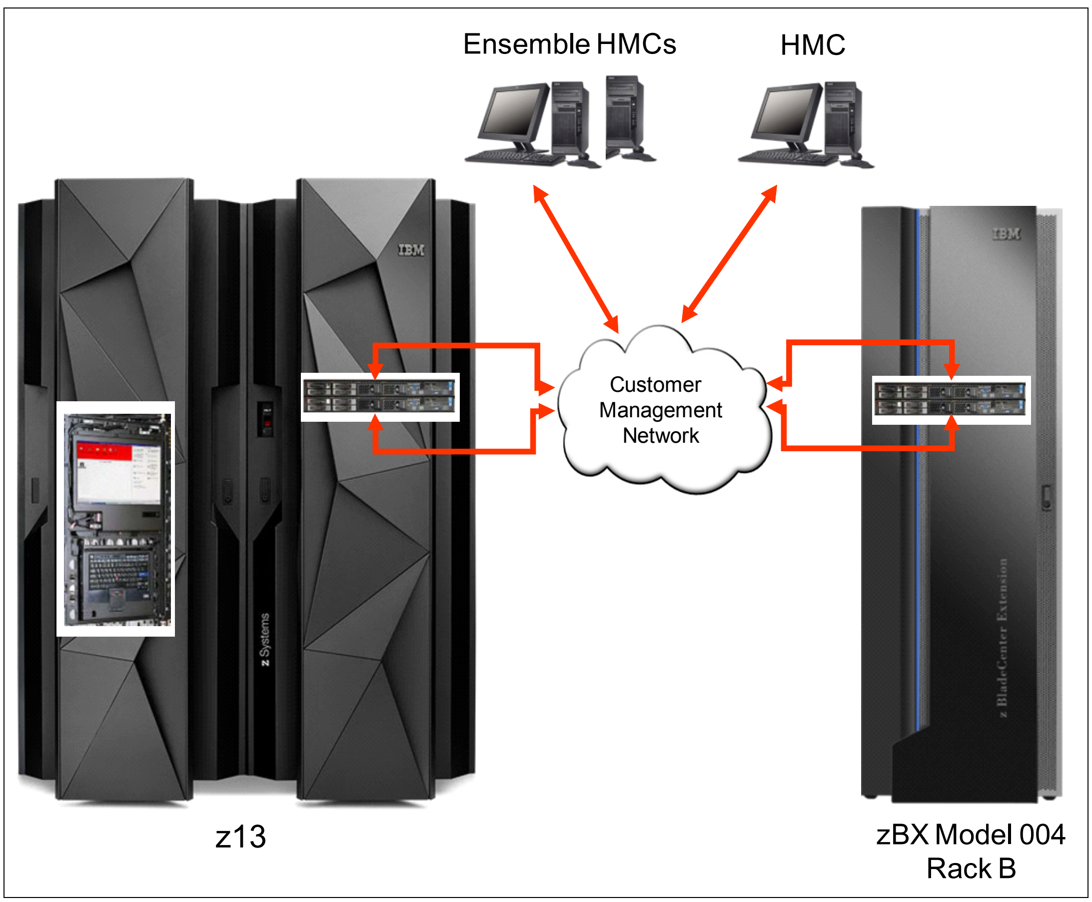

Figure 12-7 The management network

The maximum distance for the HMC Network 1000BASE-T Ethernet is 100 meters (328 feet).

Similar to the CPC there are two reserved client networks to be connected to the zBX's SEs. Customer can decide to use both networks to manage the ensemble, or use for instance Network 1 to manage the ensemble and Network 2 to provide a communication path for remote HMCs within the client intranet, and to provide the necessary broadband IBM Remote Support Facility.

An HMC that is neither the primary nor the alternate ensemble HMC can access the zBX, for instance to do single object operations, or transmit service data. However, there is no ensemble related functionality supported.

> **Note:** You can no longer order the Ethernet switches (FC 0070) that are required by a client network. Plan to use existing supported switches or acquire more switches separately to implement the required HMC LAN connectivity.

12.2.2 Intranode management network (INMN)

The INMN is one of the ensemble's two private and secure internal networks. The INMN, referred to as a *private service network*, enables the SE to communicate with the IBM blade hypervisors and optimizer blades that it controls. The HMC communicates with the SE in each zBX over a client-provided management network.

The INMN is a physically isolated 1000BASE-T Ethernet LAN, operating at 1 Gbps, that connects all resources inside the zBX for management purposes such as virtual server, performance, network virtualization, energy, and storage management. It is prewired, internally switched, configured, and managed with full redundancy for high availability.

- The INMN starts at zBX SEs and extends over INMN switches in the zBX to the BladeCenter chassis.
- The INMN switch has 48 VLAN-capable ports (counting for port 0 to 47).
- The INMN provides connectivity to the Advanced Management Module (AMM) and Ethernet Switch Module (ESM) 01/02 in each BladeCenter chassis.
- The INMN is used by the zBX's SE and HMC to run the following tasks:
  - zBX hardware monitoring and management
  - Blade management
  - Hypervisor management
  - Virtual server provisioning and management
  - Virtual network management

Figure 12-8 shows the minimum components cabled into the INMN. Some are connected through real physical network cables; others are through internal connections through midplane, or logical network connections.

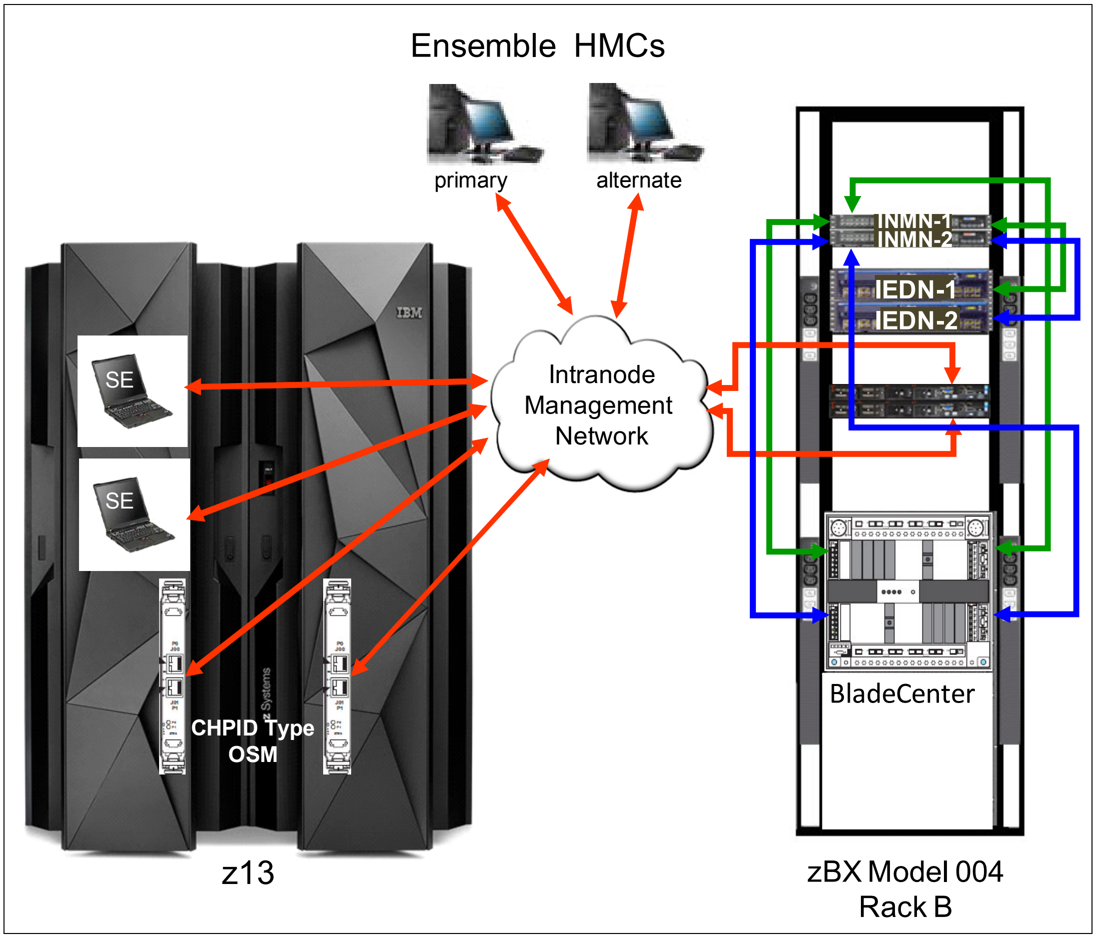

Figure 12-8 Intranode Management Network IBM z13 to zBX

The INMN uses these components:

- Primary and alternate HMCs
- Primary and alternate SEs inside the zBX B-Frame
- Two redundant CHPIDs type OSM from two separate OSA-Express5S 1000BASE-T features or OSA-Express4S 1000BASE-T (carry-forward only)
- Two redundant INMN switches that are in zBX Rack B
- Two IEDN switches cabled into the INMN switches to allow management of these switches over the INMN
- ESM01/02 and AMM in the BladeCenter connected to the INMN switches to allow management functions

OSA-Express5S 1000BASE-T feature (FC 0417)

The OSA-Express5S 1000BASE-T feature is the most recent network adapter for the CHPID OSM, OSD, OSN, OSE, and OSC in the PCIe I/O drawer in the IBM z13. There is one CHPID per card, and two ports per CHPID (Port 0 and Port 1). If the OSA-Express5S 1000BASE-T feature is configured as CHPID type OSM, it provides the connectivity to INMN from the IBM z13 for Unified Resource Manager functions. The OSA-Express4S 1000BASE-T (FC 0408) can still be used as carry forward during a MES from a zEC12.The OSA-Express3 1000BASE-T feature (FC 3367) is not supported in the IBM z13.

Figure 12-9 shows the card and port that can be used as type OSM.

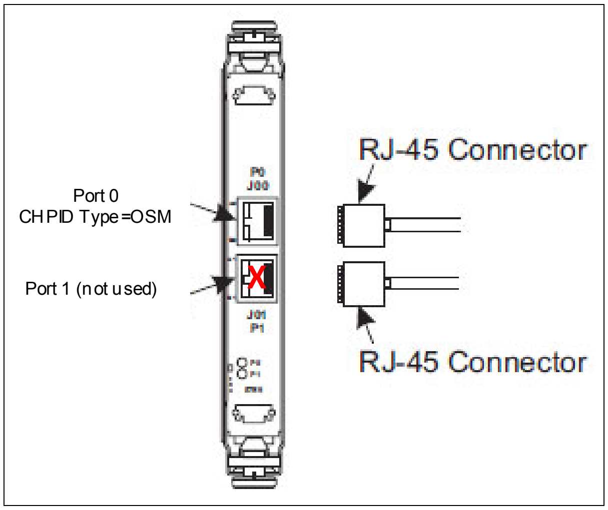

Figure 12-9 OSA-Express5S 1000BASE-T Ethernet Adapter

For using high availability, use two separate OSA-Express4S 1000BASE-T or OSA-Express5S 1000BASE-T features:

- Access to the intranode management network is available only through port 0 of any OSA-Express5S 1000BASE-T or OSA-Express4S 1000BASE-T (carry forward only) feature configured with a channel type of OSM
- The two OSM CHPIDs (for redundancy) are connected internally to port J07 of the two System Control Hubs (SCH) in the z13. They are connected using two cables provided when FC 0025 Ensemble Membership is specified on the IBM z13 initial order.
- Tell your IBM System Service Representative (SSR) which adapter will be configured as CHPID OSM. The IBM SSR will then make connections between the SCHs and the adapters.
- If FC 0025 is added to an operational IBM z13 footprint, take the OSM defined CHPIDs offline before you install the cables. The CHPIDs can then be brought online. This configuration is shown in Figure 12-10 on page 607.

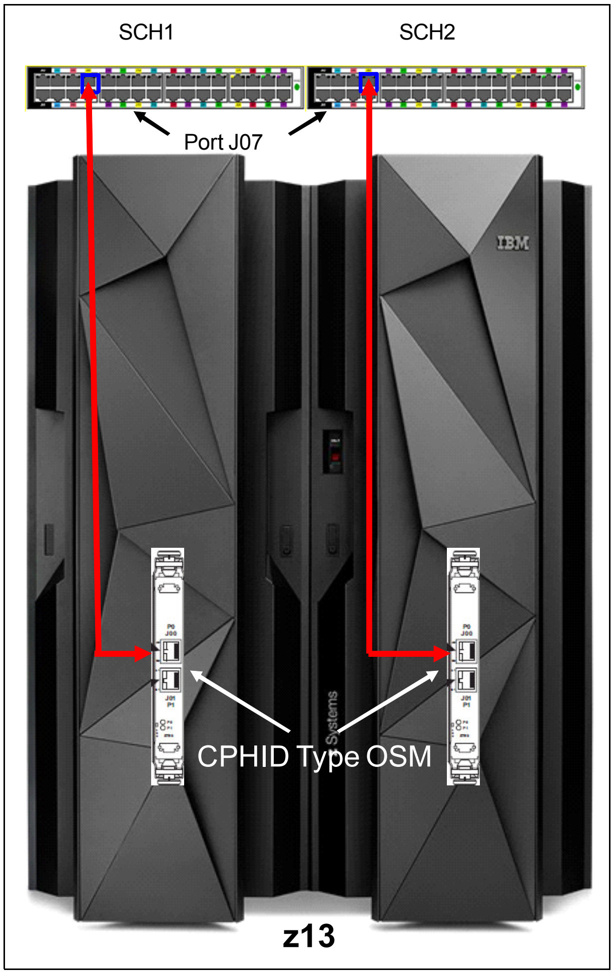

Figure 12-10 OSA-Express5S 1000BASE-T features connecting to the SCHs

> **Consideration:** OSA-Express4 1000BASE-T feature (FC 0408) can be used as a CHPID OSM in the IBM z13. However, it is carried forward only on system MES.

Intranode management network switches

The INMN switches are mounted in the rear of the zBX Rack B. They are located above the IEDN switches. The TOR switches for INMN are EIA B36P(INMN-A) and EIA E35P(INMN-B). All connections to both switches are made at the rear of the rack. Consider this information:

- Both INMN switches have 48 ports each.
- All ports operate at 1 Gbps. All cables are pre-wired and supported by the controlling SE in the zBX's B-Frame.
- All cables connecting to the INMN switches are internal to the zBX racks.

Each BladeCenter chassis has two cable paths per INMN switch for redundancy. All INMN Ethernet cables to the BladeCenters chassis and the IEDN switches are connected at installation time.

The port assignments for both INMN TOR switches are listed in Table 12-2.

Table 12-2 Port assignments for the INMN TOR switches

Ports	Description
J00-J03	Management for BladeCenters in zBX Rack-B
J04-J07	Management for BladeCenters in zBX Rack-C
J08-J11	Management for BladeCenters in zBX Rack-D
J12-J15	Management for BladeCenters in zBX Rack-E
J16-J41	Not used
J42	INMN-A to SE- A (eht1) / INMN-B to SE- A (eht2)
J43	INMN-A to SE- B (eht1) / INMN-B to SE- B (eht2)
J44-J45	INMN switch B36P (Top) to INMN switch B35P (Bottom)
J46	INMN-A to IEDN-A port J41 / INMN-B to IEDN-B port J41
J47	INMN-A to SE A (IMM)/ INMN-B to SE B (IMM)

Figure 12-11 shows how the INMN cabling looks for a configuration of two zBX racks with four BladeCenters installed. Each BladeCenter chassis has two cable paths per switch for redundancy that are cabled to the ESM and the AMM in the BladeCenter chassis.

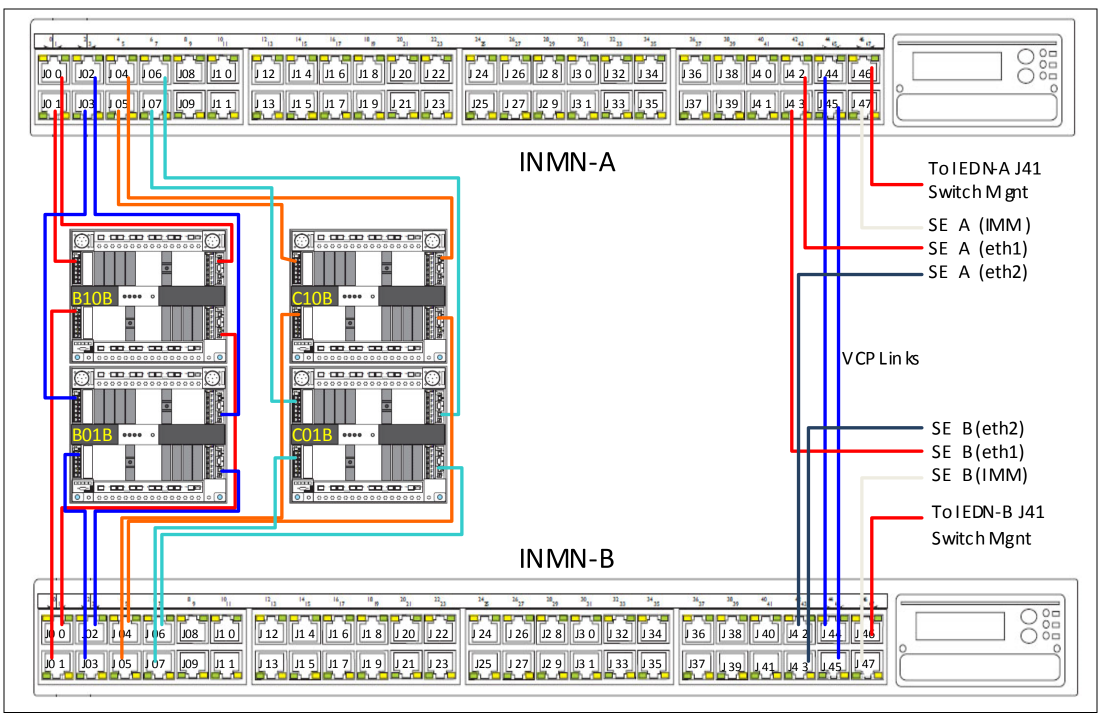

Figure 12-11 zBX INMN TOR switch wiring with four zBX BladeCenters (two rack configuration)

Remember: The two INMN switches are present only in zBX Rack B. All other BladeCenter chassis are connected to them.

For more information about setting up OSM definitions see these sections:

- 14.3.1, “Defining the OSM definitions” on page 673
- 14.4.1, “z/OS display commands for CHPID type OSM” on page 683

12.2.3 Intraensemble data network (IEDN)

An IEDN is sometimes referred to as a *private data network*. An IEDN is visible to the operating systems that run in virtual servers, and is required for application data communications within an ensemble that contains one zBX or more.

The IEDN is a 10 Gigabit Ethernet (GbE) fiber optic connection between the CPCs and the zBX. Any CPC connecting to the zBX through a IEDN requires a port on two separate OSA-Express5S 10 GbE, or two separate OSA-Express4S features (carry forward only). The zBX Model 004 can support up to eight separate IEDN connections to other CPCs, and up to eight IEDN connections to other zBXs.

In zBX Model 004, an IEDN provides a potential of 20 GbE bandwidth through link aggregation. There are two 10 GbE cables between HSS modules 07/09 and the IEDN switches.

All the fiber optic cables that connect the zBX to the IBM z13 through the IEDN must be provided by the client:

- Use two ports to provide a redundant configuration for failover purposes in case one link fails.
- For high availability, spread the two ports across two different OSA-Express4S 10 GbE, or OSA-Express5S 10 GbE features within the same CPC and in different I/O domains (in the PCIe I/O drawer), if possible.
- The IEDN is used exclusively for point-to-point data traffic. This means that no routers or switches are allowed between the IBM z13 and the TOR switches inside the zBX.
- The client is responsible for all external fiber optic cables that connect the IBM z13 to the zBX IEDN switch.

Figure 12-12 on page 610 shows the basic components in the IEDN. Some are connected through external network cables, and others through internal connections via midplane or logical network connections:

- Two redundant IEDN switches are in zBX Rack B.
- Two OSA-Express-5S 10 GbE features (CHPID type OSX), either FC 0415(LR) or FC 0416(SR). Also two OSA-Express-4S 10 GbE features (CHPID type OSX), either FC 0406(LR) or FC 0407(SR) can be used as a carry forward during MES from a zEC12.
- Two High Speed Switch (HSS) modules 07/09 in the BladeCenter chassis.

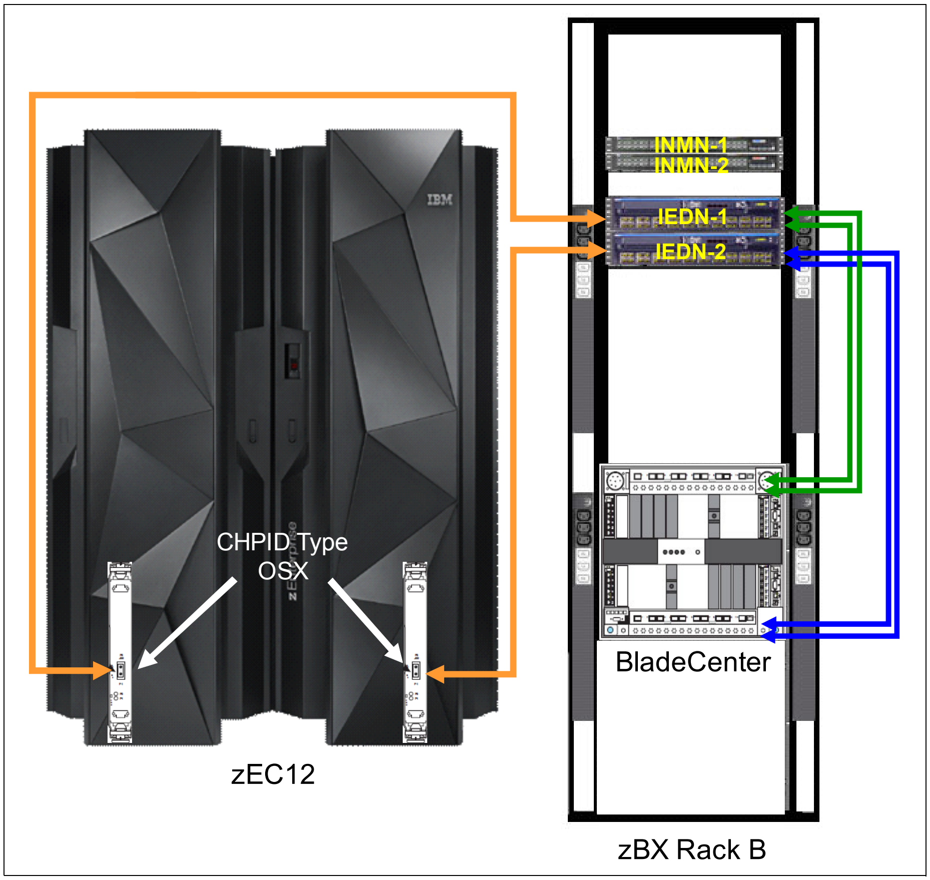

Figure 12-12   Intraensemble Data Network connections

Also, the IEDN can support an internal LAN segment for the HiperSockets network. Each CPC can have one IQD CHPID with IEDN access defined to enable IEDN HiperSockets communication to the virtual servers. This configuration enables the HiperSockets network to extend outside the CPC to the entire ensemble.

OSA-Express5S 10 GbE features (FC 0415/0416)

FC 0415 (LR) and 0416 (SR) are the most recent network features for the CHPID OSX. These features provide the connectivity and access control for IEDN to the IBM z13.The CHPID type OSX connects from the z13 to the IEDN TOR in Rack B.

These features can also be configured as CHPID type OSD when not part of the IEDN network.

The minimum requirement for CHPID type OSX is two OSA-Express5S/4S 10 GbE adapters in separate domains in each IBM z13. These adapters have only one port as shown in Figure 12-13 on page 611.

For information about setting up OSA-Express for zBX (OSX) definitions, see these sections:

- 14.3.2, “Defining the OSX definitions” on page 678
- 14.4.2, “z/OS display commands for CHPID type OSX” on page 684

> **Note:** Legacy OSA-Express3 cards are not supported in the z13. This applies to OSA-Express3 1000BASE-T feature (3367) and OSA-Express3 10 GbE features (3370 and 3371).

Figure 12-13 shows the adapter, the LC cable connector, and their specifications.

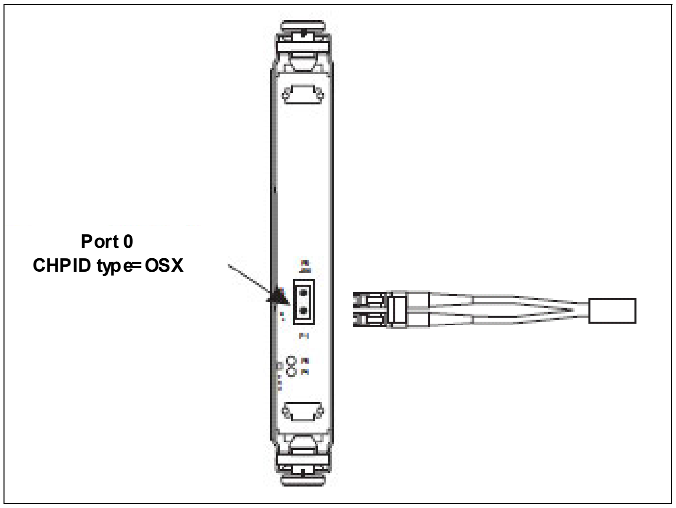

Figure 12-13 OSA-Express5S 10 GbE feature

Intraensemble data network (IEDN) switches

The 2U IEDN switches are mounted from the rear of the zBX Rack B and located below the INMN TOR switches. The switches are IEDN-A (EIA 32P upper) and IEDN-B (EIA 30P lower). All fiber optic connections are at the rear of the zBX Rack B.

Figure 12-14 shows the IEDN switch port designations. All internal fiber optic connections from the BladeCenter chassis to the IEDN switches in Rack B are pre-plugged. If there are more racks, the fiber connections from the BladeCenter chassis within those racks are plugged at installation time. These cables are included with the features ordered.

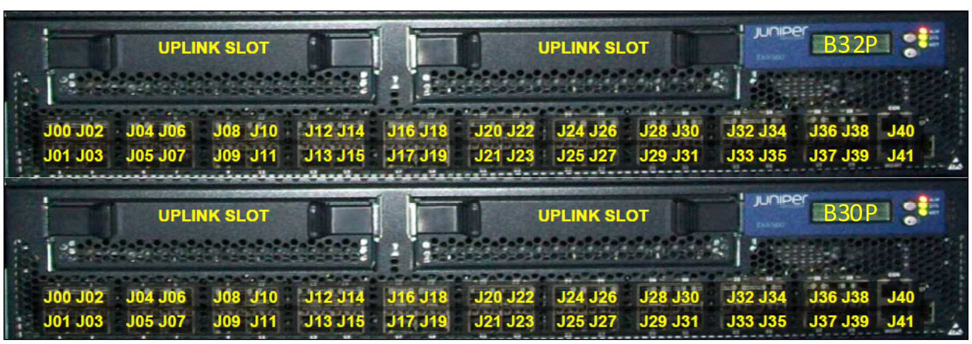

Figure 12-14 Intraensemble Data Network switches

The port assignments for both IEDN TOR switches are listed in Table 12-3 on page 612.

Table 12-3 Port assignments for the IEDN TOR switches

Ports	Description
J00 - J07	SFP+ reserved for zEC12(OSX) IEDN connections
J08 - J23	DAC reserved for BladeCenter and SM07/SM09 IEDN connections
J24 - J30	SFP+ reserved for zBX-to-zBX IEDN connections
J31 - J37	SFP+ or SFP reserved for client managed network connections
J38 - J39	1 Meter DAC for Switch to Switch VC
J40	RJ-45 console port
J41	RJ-45 IEDN switch management port to INMN switch port 46

The IEDN TOR switches use Small Form Factor Pluggable (SFP) optics for the external connections, and Direct Attach Cables (DACs) for internal connections, as shown in Figure 12-15.

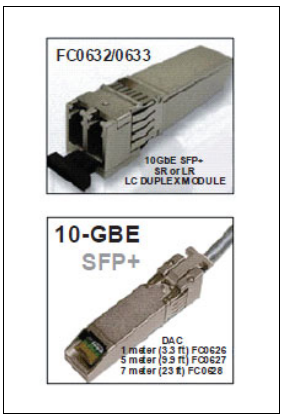

Figure 12-15 IEDN SFPs and DAC cables

- The IEDN is a 10 GbE flat data link (Layer 2) network.
- Communications Server for z/OS requires minimal configuration:
  - IPv4 or IPv6 addressing.
  - VLAN IDs must be configured and must match the zManager configuration.
- The z/VM virtual switch types[2] provide IEDN access:
  - The uplink can be a virtual network interface card (vNIC).
  - Ensemble membership conveys an Ensemble Universally Unique Identifier (UUID) and media access control (MAC) prefix.
- The IEDN network definition is done from the primary ensemble HMC by using the Manage Virtual Network task.

Requirement: Connections from client's external network to the ensemble via the IEDN switches must be a routed connection and not a switched connection. The IEDN switch is configured to expect Layer 3 routed connectivity and so is incompatible with typical Layer 2 switching protocols.

[2] Only z/VM 6.2 supports the IEDN. The z13 does not support z/VM integrated systems management. Therefore, z/VM 6.2 does not support the IEDN on z13.

The IEDN switch BladeCenter cables and IEDN-A to IEDN-B cables are internal and are provided with the zBX Direct Attach Cables (Figure 12-6 on page 602). Also shown are the Ethernet port designations for the switch management cables, which connect to the INMN switches.

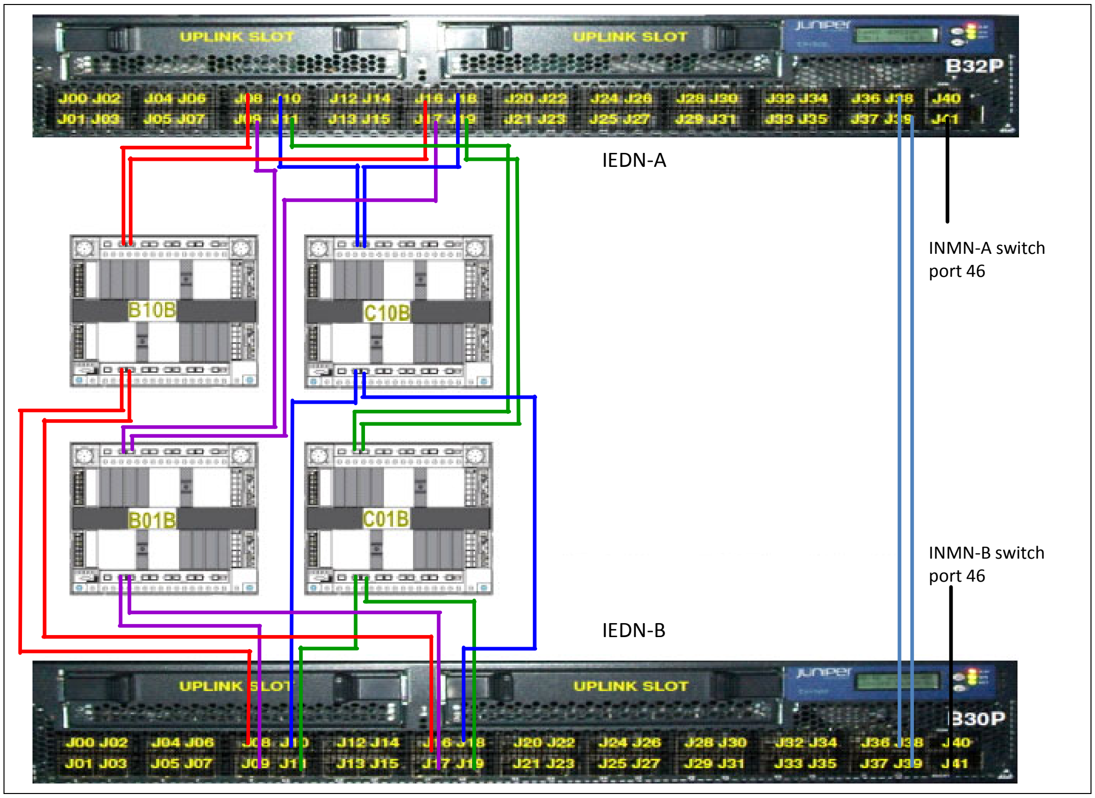

Figure 12-16 zBX IEDN internal switch wiring with four zBX BladeCenter chassis

zBX to zBX IEDN connection

If more than one zBX Rack B is in an ensemble, both must be connected to each other through IEDN switches. Here are the considerations when connecting zBX to zBX in IEDN:

- The client must provide all cables for zBX to zBX connections.
- There are two masters (node1 and node8) in an ensemble that consist of more than one zBX Rack B.
- Other than the master zBX, adding/removing nodes is concurrent.
- The replacement of all nodes is concurrent, even for the master node.
- The first zBX is a node1 when the ensemble is created or added to an existing ensemble. The second zBX is added to an existing ensemble as a node8. The zBXs are then added from a node2 to a node7. For more information, see these publications:
  - *Installation Manual for Physical Planning 2964*, GC28-6938
  - *Installation Manual for Physical Planning 2458-004*, GC27-2630

- In Figure 12-17, if you remove node 2, node3 becomes node2. Therefore, the cabling for node 3 must be changed to match that of node 2.

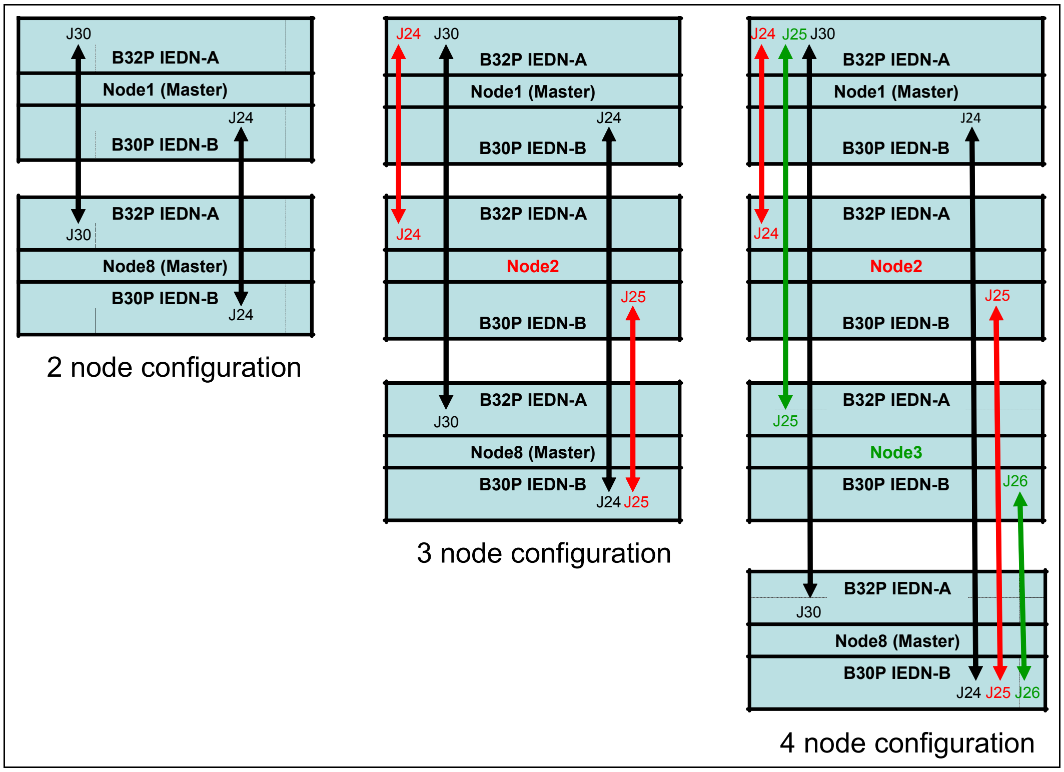

Figure 12-17 zBX to zBX IEDN connection

For information about setting up OSX definitions see these sections:

- 14.3.2, “Defining the OSX definitions” on page 678
- 14.4.2, “z/OS display commands for CHPID type OSX” on page 684

12.2.4 Fibre Channel (FC) attached storage

Fibre Channel (FC) connections can be established between the zBX and a SAN environment through two 20-port 8 Gb FC switch modules (SM03 and SM04) in each BladeCenter chassis. All switch modules run in pass-through mode, and have 4 internal ports and 6 external ports.

Each IBM blade in the BladeCenter chassis has an expansion adapter, a two-port adapter that is connected into the zBX with SM03 and SM04 switches. This internal topology is already set up when the zBX is delivered.

Figure 12-18 shows the FC connection between a BladeCenter chassis and external storage.

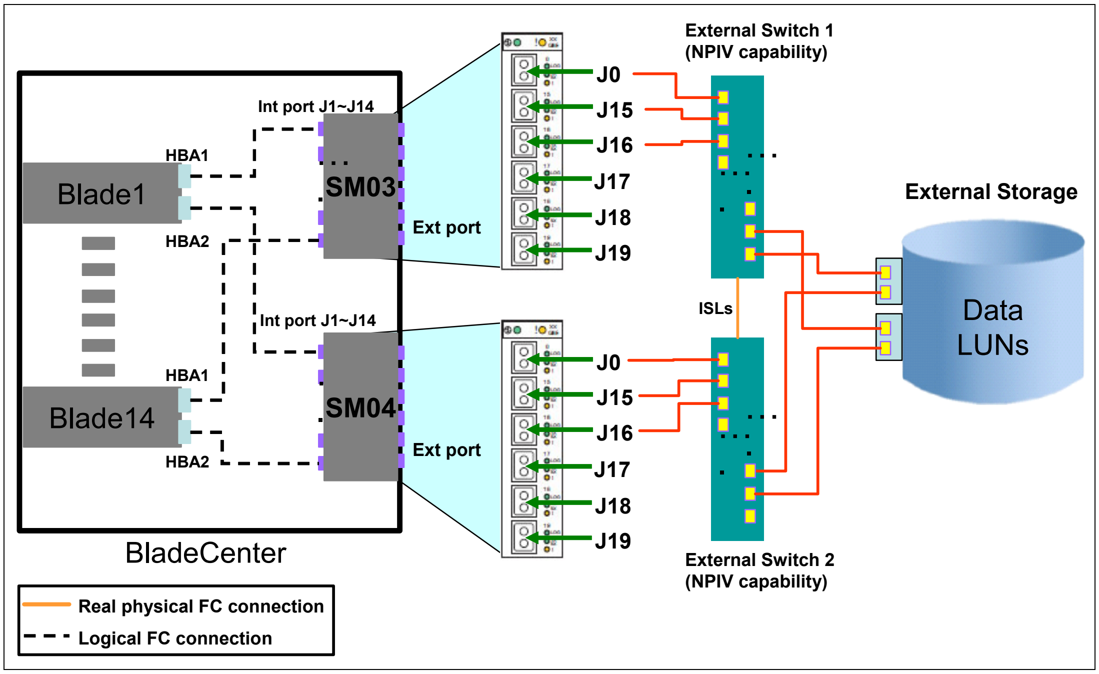

Figure 12-18 BladeCenter storage topology and Switch Modules03 and 04

Some implications exist for the SAN configuration:

- All components of the SAN are client-owned, client-supplied, and client-managed.
- Each switch module in the BladeCenter chassis supports up to six connections to the SAN fabric. Only SM03 and SM04 are used to make connections to disk storage. All other modules are reserved for other purposes.
- Each switch module in the BladeCenter chassis maps the blades to one of six external ports when a blade runs Fabric Login (when the blade is activated). The switch modules usually select the external port with the fewest blades mapped when a new blade runs a Fabric Login.
- The first SAN switch must support N-Port ID Virtualization (NPIV) mode and operate at 2, 4, or 8 Gb per second. A speed of 1 Gbps is not supported. Review the interoperability matrix for the most recent details:

  http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss
- The configuration does not add an extra hop to your SAN. Some SAN switches have a limit on the number of hops allowed between the host port and the target port.
- Separate SAN switches can be used for high availability but is not a requirement.

The maximum distance depends on the speed, fiber optic cable diameter, and signal frequency, as shown in Table 12-4.

Table 12-4 FC SAN fiber optic cable types: Maximum lengths

Fiber optic cable types	8 Gbps	4 Gbps	2 Gbps
50 micron (2000 MHz-km)	150 meters 492 feet	380 meters 1247 feet	500 meters 1640 feet
50 micron (500 MHz-km)	50 meters 164 feet	150 meters 492 feet	300 meters 984 feet
62.5 micron (200 MHz-km)	21 meters 69 feet	70 meters 230 feet	150 meters 492 feet

12.3 Upgrade considerations

An upgrade of the zBX Model 002 or Model 003 to Model 004 becomes necessary if the owning zEC12, or z196, will be replaced by an IBM z13. You can decide to request the upgrade to have the zBX decoupled from the owning CPC and upgrade the CPC at a later time.

This section lists hardware installation considerations when you upgrade to a zBX Model 004. It becomes a stand-alone node in the ensemble and all connections which existed in the zEC12 ensemble, except the INMN connection from CPC the zBX, remains in place.

Ensemble considerations

Consider the following information related to ensembles:

- A Systems Assurance Technical Delivery Assessment (TDA) meeting is scheduled to support the installation. This is the official planning vehicle for a smooth installation.
- Two HMCs are required to manage a zBX. The machine type of the primary HMC and alternate ensemble HMC must be identical. Both must at least be FC 0091 (7327). Verify this by checking the label on the HMC tower (FC 0091/0092) or HMC rackmounted chassis (FC 0094).
- HMC feature FC 0091 must have at least 16 GB RAM.
- Both the primary and alternate Ensemble HMC must be connected on the same Ethernet subnet.
- Most recent HMC feature is the FC 0094 (7914).
- HMC application version 2.13.0 is required for an IBM z13 ensemble.
- Consider upgrading the ensemble HMCs a few days before the zBX upgrade.
- Software Preventive Service Planning (PSP) buckets for the zBX must be reviewed. The PSP buckets are available on the IBM Resource Link under the **Fixes** → **PSP** menu. Search for `2458DEVICE`.
- Review the System Storage Interoperation Center (SSIC) website for prerequisites in the Open Systems environment:

  http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss

zBX considerations

Consider the following information related to zBX:

- If the zBX Model 004 is relocated, environmental requirements including floor cutouts, perforated tiles positioned at the front of the racks, client supplied power for the racks, and installation positioning must be planned and prepared in advance.
- The relocation of the zBX can also require a replacement of the existing cables in the ensemble. This can include the replacement of cables for these instances (see Figure 12-20 on page 622):
  - the management network connecting to the SEs in the zBX
  - OSX connections
  - zBX to zBX connections
  - zBX connections to customer network
  - SAN connections to a SAN switch
- All fiber optic cables that support external attachment to the zBX IEDN network are the client's responsibility.
- If the B-Frame of the zBX Model 003 consists only of a single BladeCenter chassis, the upgrade to Model 004 adds two power distribution units (PDUs) to the frame. The must provide two more connections to wallpower.
- The zBX Model 004 cannot be purchased from IBM as a new device. Only a existing zBX Model 002 or Model 003 can be upgraded to a zBX Model 004.
- No additional racks or BladeCenter chassis can be added during MES.
- ISAOPT Blades are not supported in a zBX Model 004.
- The zBX upgrade is disruptive for all blade operations.
- An upgrade from a zBX Model 002 to zBX Model 004 is possible and includes additional physical rework (reworking and extending the zBX internal IEDN Network).
- Blades from a zBX Model 002/003, except IOSAOPT Blades, can be carried forward to the zBX Model 004.
- If customer is upgrading to zBX Model 004 and moves existing blades by Exporting zBX data from the owning zEC12, or z196 (Donor CPC), any DP blades has to be removed before. DP blades can be installed afterwards when the MES has been completed.
- During MES customer is able to purchase additional entitlement up to the maximum capacity of the existing zBX hardware (up to the amount of free slots in the existing BladeCenter chassis).
- When the entitlement has been purchased and installed, customer has to purchase the zBX supported Blades to be added from a vendor other than IBM.
- Because the zBX Model 004 becomes an independent node in the IBM z13 ensemble, the customer must provide network-related information for SE configurations, including these items:
  - Primary TCP/IP and alternate TCP/IP Address, plus TCP/IP Network Mask
  - The IP of a Domain Name Server
  - Domain Suffix Search Order
  - The Gateway Address

z13 considerations

Consider the following information related to IBM z13:

- To participate in an ensemble, any CPC must meet these requirements:
  - Installed Ensemble Membership Flag (FC 0025)
  - Installed QoS Feature (FC 0019 level Manage Firmware Suite or FC 0020 level *Autotmate/Advanced Management Suite*). All nodes in the same ensemble must have the same QoS Feature installed.
  - All CPCs becoming a member in an ensemble must have at least two OSA-Express5S 1000BASE-T or two OSA-Express4S 1000BASE-T ports (carry forward only) for INMN connections installed. IBM supplies the cables to connect the OSA-Express 1000BASE-T adapters to both Bulkpower Hubs (BPHs), if the CPC is a System z196 or a System zEC12, or to both System Control Hubs, if the CPC is an IBM z13. These features must be defined as CHPID type OSM.

    For more information about setting up OSM definitions, see these sections:

    - 14.3.1, “Defining the OSM definitions” on page 673
    - 14.4.1, “z/OS display commands for CHPID type OSM” on page 683
  - All CPCs connecting to the zBX over the IEDN must have at least two OSA-Express5S 10 GbE ports or OSA-Express4S 10 GbE ports (carry forward only) installed. The client is responsible for providing the OSA cables that connect the IEDN switches and the CPC. These features must be defined as CHPID type OSX.

    For information about setting up OSX definitions, see these sections:

    - 14.3.2, “Defining the OSX definitions” on page 678
    - 14.4.2, “z/OS display commands for CHPID type OSX” on page 684
- Ensure that you have the correct SFP/SFP+ module feature code and DAC feature code for the IEDN switches. The ports on the IEDN switches are segmented by purpose:
  - Connection to client OSD ports (J31-J37):
    - 10 GbE LR SFP+ module (FC 0603)
    - 10 GbE SR SFP+ module (FC 0633)
    - 1 GbE 1000BASE-LX SFP module (FC0634)
    - 1 GbE 1000BASE-SX SFP module (FC0635)
  - Connection to BladeCenter SM07/SM09 (J8-J23):
    - 10 GbE 1m DAC (FC 0626)
    - 10 GbE 3m DAC (FC 0027)
    - 10 GbE 7m DAC (FC 0028)
- All fiber optic cables that support external attachment to the zBX IEDN network are the client’s responsibility. These fiber optic cables can connect from zBX to zBX, or from a CPC to a zBX. All necessary Fibre Channel cables for the installation are ordered or planned for the following connections:
  - OSA-Express5S/4S 10GbE ports to the IEDN switches
  - zBX IEDN to zBX IEDN

12.4 Upgrade overview

This section provides a high-level overview of the zBX upgrade from Model 003 to Model 004. The sequence might change over time. Contact your IBM System Service Representative (SSR) for information about the most recent changes.

The upgrade to Model 004 installs these components into Rack B:

- 2 X Support Element Server (SEs)
- 2 X Trays consisting of Keyboard and Display (KVM)
- 2 X Power Distribution Units (PDUs) if Rack B has only one BladeCenter chassis
- 6 X 1 Gb Ethernet Cables (Cat 6)
- 1 X 1U Filler plate

The zBX upgrade involves these tasks:

- Preparing the upgrade:
  - Gathering all configuration information (management network informations).
  - Planning all actions for cable or environmental rework, if the zBX is relocated.
  - Estimating installation times. The zBX upgrade takes approximately eight hours depending on the configuration of the zBX system. Adding blades during the upgrade takes additional time depending on the amount of blades to be added.
- Replacing or upgrading the ensemble HMCs

Upgrade Path A (assuming that zBX Model 004 will stay in the same ensemble) involves these tasks:

- Exporting zBX data from the owning zEC12 or z196 (Donor CPC)
- Detaching or removing zBX for the owning CPC (including deactivation of the blades)
- Powering off the zBX
- On the owning CPC, converting zBX Feature to Ensemble without zBX

Upgrade Path B (assuming that zBX Model 004 will move into a new or different ensemble) involves these tasks:

- Deactivating virtual servers
- Deleting virtual servers and remove storage resources
- Removing blade entitlement
- Detaching/removing the zBX for from the owning CPC
- Powering off the zBX
- On the owning CPC, removing zBX Hardware from CPCs configuration
- On the owning CPC, converting zBX Feature to Ensemble without zBX (assuming that the CPC will remain in the ensemble)

Resuming installation (valid for Upgrade Path A and B) involves these tasks:

- Adding two extra PDUs to Rack B (only if the zBX has only one BladeCenter chassis installed).
- Moving cables for zBX Model 002 to Model 004 conversion. This does not apply if a zBX Model 003 is upgraded.
- Installing the SEs, the 1U filler plate, and the trays with display and keyboard.
- Installing the Ethernet cables to connect zBX SEs to the INMN.
- Powering on the zBX.

Following Upgrade Path A involves this task:

- Restoring previously saved zBX Data

Following Upgrade Path B involves this task

- Installing the zBX Model 004 into the new ensemble

12.5 Configuration examples

This section presents configuration examples that contain a zBX, and the necessary connectivity requirements for operation. Follow-on configuration diagrams build on the first configuration and only extra requirements are noted.

12.5.1 Basic Ensemble configuration with zBX

Figure 12-19 shows one IBM z13 (CPC1), a zBX Mod 004, switches, and FC attached storage.

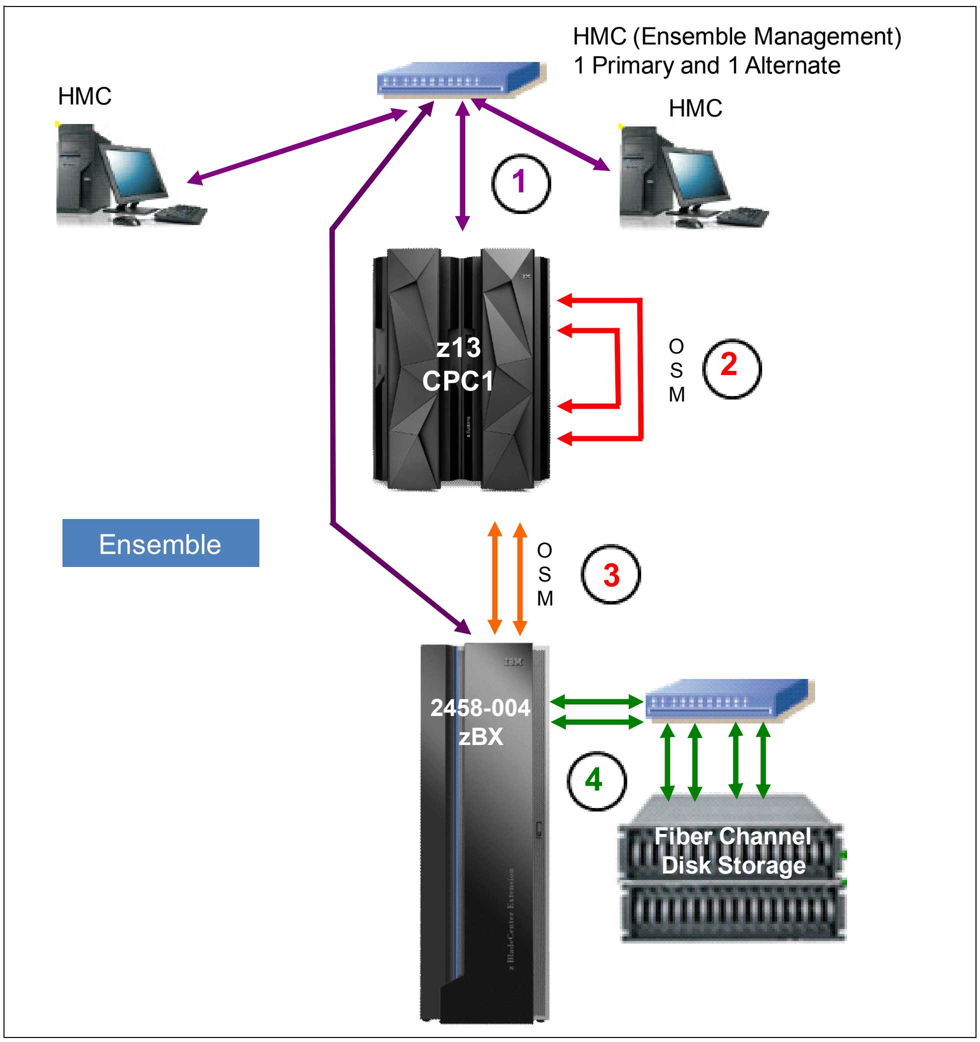

Figure 12-19 Single CPC with zBX

The following list indicates the necessary connections, required cabling, and necessary components within the ensemble:

- Client-provided management network:
  - Two ensemble HMCs (primary and alternate)
  - IBM supplied Ethernet RJ-45 cables (50 feet)
  - 1 Gb compatible client-provided network switches
  - The switch connection to the SEs in the IBM z13
  - The switch connection to the SEs in the zBx Model 004
- INMN:
  - Two redundant CHPID type OSM from two different OSA-Express5S 1000BASE-T, or OSA-Express4S 1000BASE-T (carry forward only) features.
  - IBM supplies two Ethernet Category 6 cables (for redundancy) from the OSM CHPIDs to IBM z13 SCH in location A31BSCH1-J07/A31BSCH2-J07. This is a internal connection in the IBM z13 that is supplied with FC 0025.
- IEDN:
  - Two redundant ports from two different OSA-Express5S 10 GbE features (SR or LR), or OSA-Express4S 10 GbE (carry forward only) features (SR or LR) configured as CHPID type OSX.
  - The client supplies the fiber optic cables (single mode or multi-mode).
- FC attached disk storage:
  - The client supplies all fiber optic cables (single mode or multi-mode) from the zBX to the attached disk storage.
  - The client supplies switches/routers if necessary, and is responsible for the management of the FC attached disk storage.

Note: Ordering the Ethernet switches (FC 0070) required by HMCs is no longer possible. Plan to use existing supported switches or acquire more switches separately to implement the required HMC LAN connectivity.

12.5.2 Ensemble configuration with two zBXs and two CPCs

The ensemble shown in Figure 12-20 on page 622 consists of a zEC12 (CPC1) with an attached zBX Model 003, an IBM z13 (CPC2), a stand-alone zBX Model 004, and FC attached disk storage to each zBX.

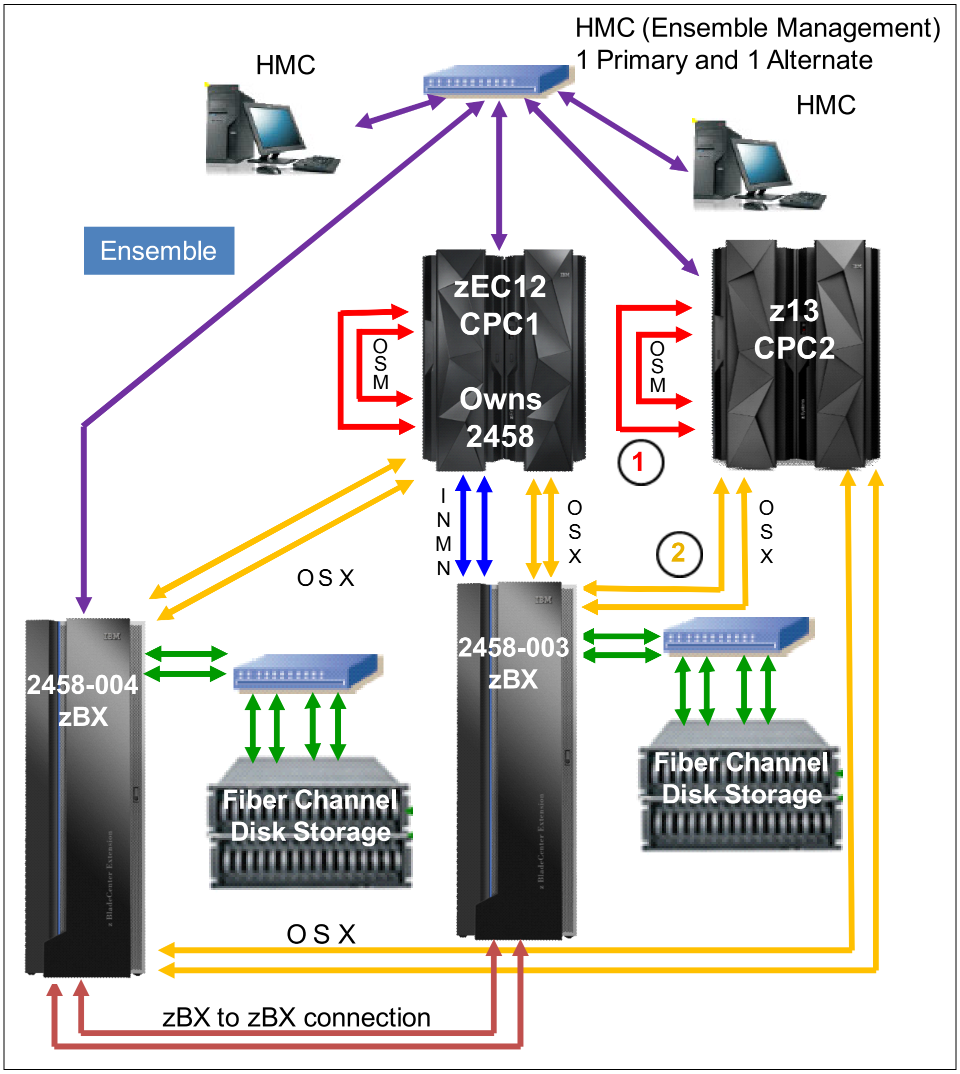

Figure 12-20 Dual CPC Ensemble with zBX

This configuration has these characteristics:

- INMN:
  - Two redundant CHPIDs type OSM (for each CPC) using OSA-Express5S or OSA-Express4S 1000BASE-T (carry forward only) features.
  - IBM supplies two Ethernet Cat6 cables from the OSM CHPIDs to IBM z13 System Control Hubs (SCHs) in location A31BSCH1-J07/A31BSCH2-J07.
  - IBM supplies two Ethernet Cat6 cables from the OSM CHPIDs to IBM z13 System Control Hubs (BPHs) in the zEC12.
- IEDN:
  - Two redundant ports from two separate OSA-Express4S/5S 10 GbE features (SR or LR) are configured as CHPID type OSX. (four OSX PCHIDs need for each CPC).
  - The client supplies the fiber optic cables (single mode or multi-mode).

12.5.3 Maximum mixed CPC/zBX Ensemble configuration

Figure 12-21 shows the maximum connected CPCs (eight) to a zBX using the IEDN. Three z196 CPCs are connected by OSA-Express3 ports (CHPID type OSD). They are either direct-connected, or connected through client-supplied switches.

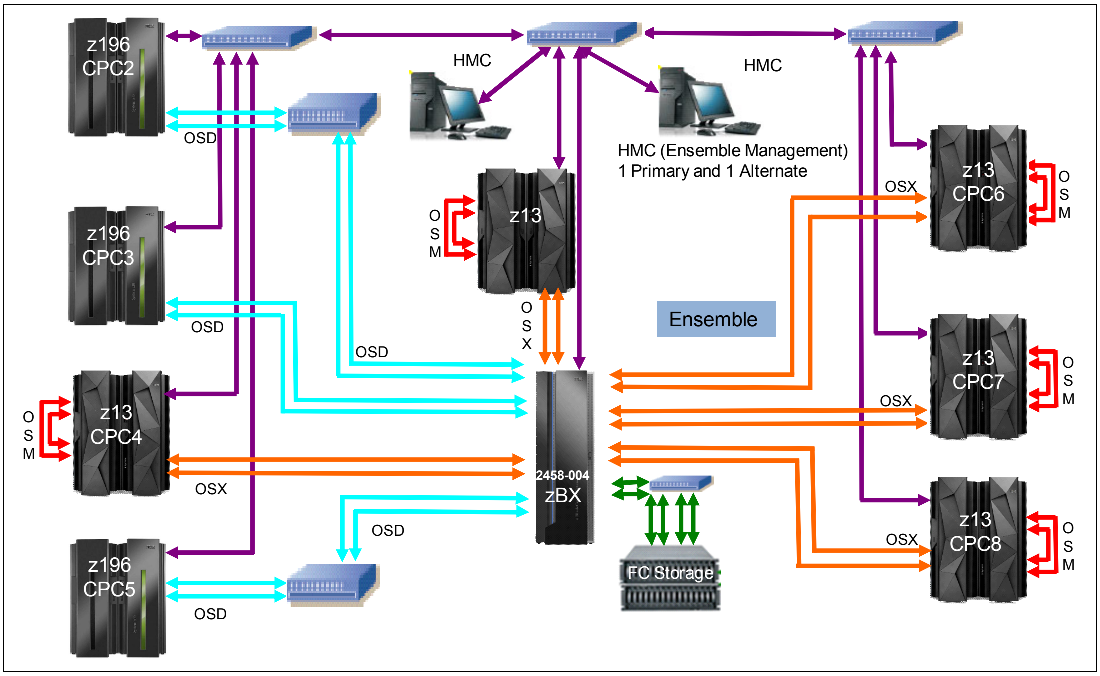

Figure 12-21 Maximum CPC configuration (eight CPCs connected over an IEDN)

The z196s are not part of the ensemble, but they are allowed to use the zBX applications. More IBM z13 CPCs are added and connected to the zBX over the OSA-Express4S/5S 10 GbE features (SR or LR) configured as CHPID type OSX.

12.6 Ensembles and Unified Resource Manager

An *ensemble* is a platform systems management domain that consists of one or more CPCs (at least System z196) with or without zBX. Through the ensemble, z Systems and other server resources are effectively integrated into a single platform. The ensemble provides an integrated way to manage virtual server resources and the workloads that can be deployed on those resources.

12.6.1 Ensemble creation

The ensemble starts with a pair of HMCs that are assigned to an ensemble identity. In an ensemble, all nodes (CPCs and zBXs) are then added to the ensemble through an explicit action at the HMC. A new IBM z13 must be ordered with FC 0025 (Ensemble Membership) that is associated with HMC enablement. For more information about HMC considerations in an ensemble, see 3.2.11, “Connectivity for the HMC and the SE with ensemble enabled” on page 91.

Ensemble membership requires that the IBM z13 has two CHPID type OSMs with a connection to port 7 of the SCHs over two Ethernet cables. FC 0025 provides these cables and ensures that the necessary feature code for OSA cards is present. The connection for these cables is described in 12.2.2, "Intranode management network (INMN)" on page 604.

The HMC with FC 0025 has a new task, **Create Ensemble**, under Console Actions Work Area (Figure 12-22). When the ensemble is created, the **Ensemble** icon is added to the Views area of the HMC.

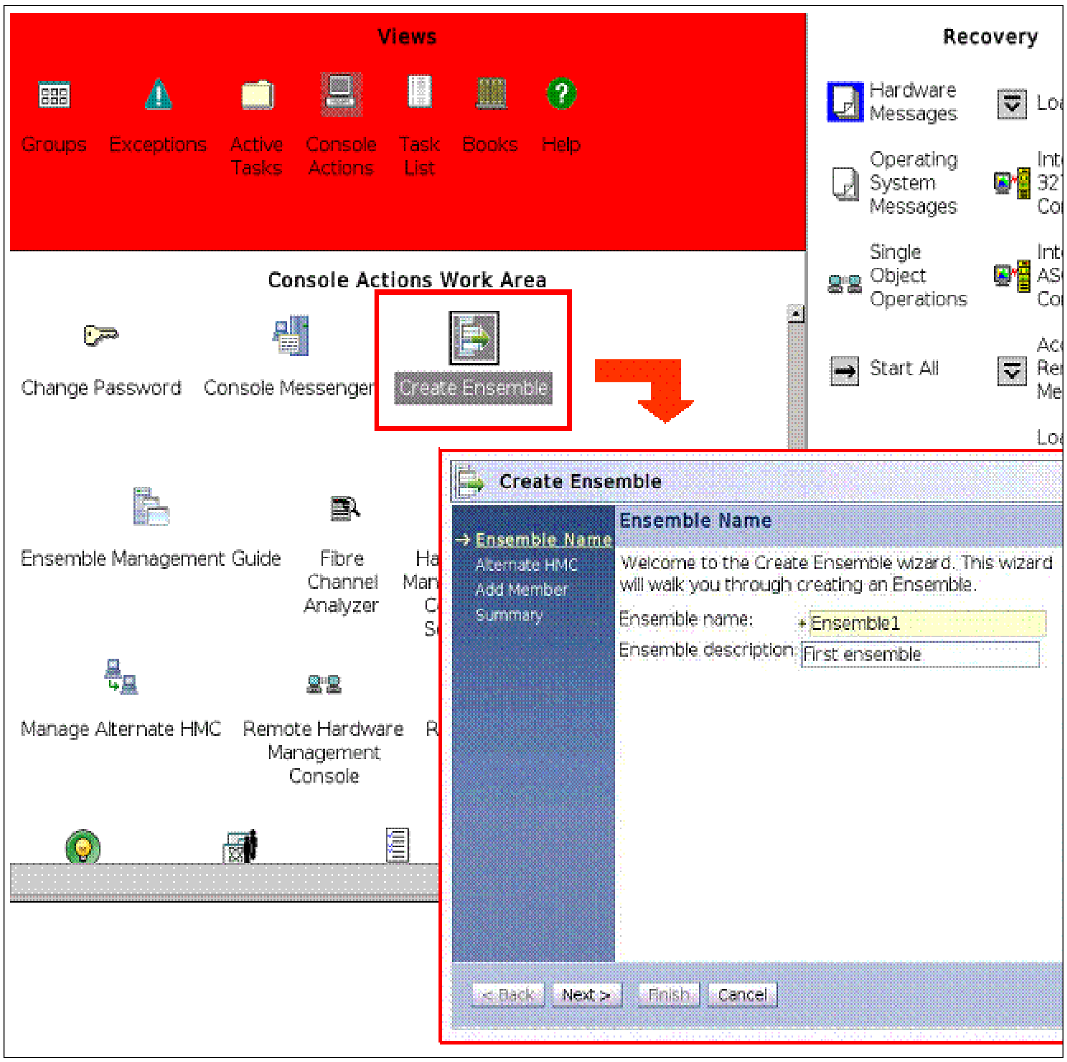

Figure 12-22 Ensemble creation task

Before you create the ensemble on the primary HMC, you must have a candidate HMC to be used as the alternate HMC. Figure 12-23 shows a Manage Alternate HMC task in the Console Actions Work Area. This task is used to define a candidate HMC that will be used as the alternate HMC.

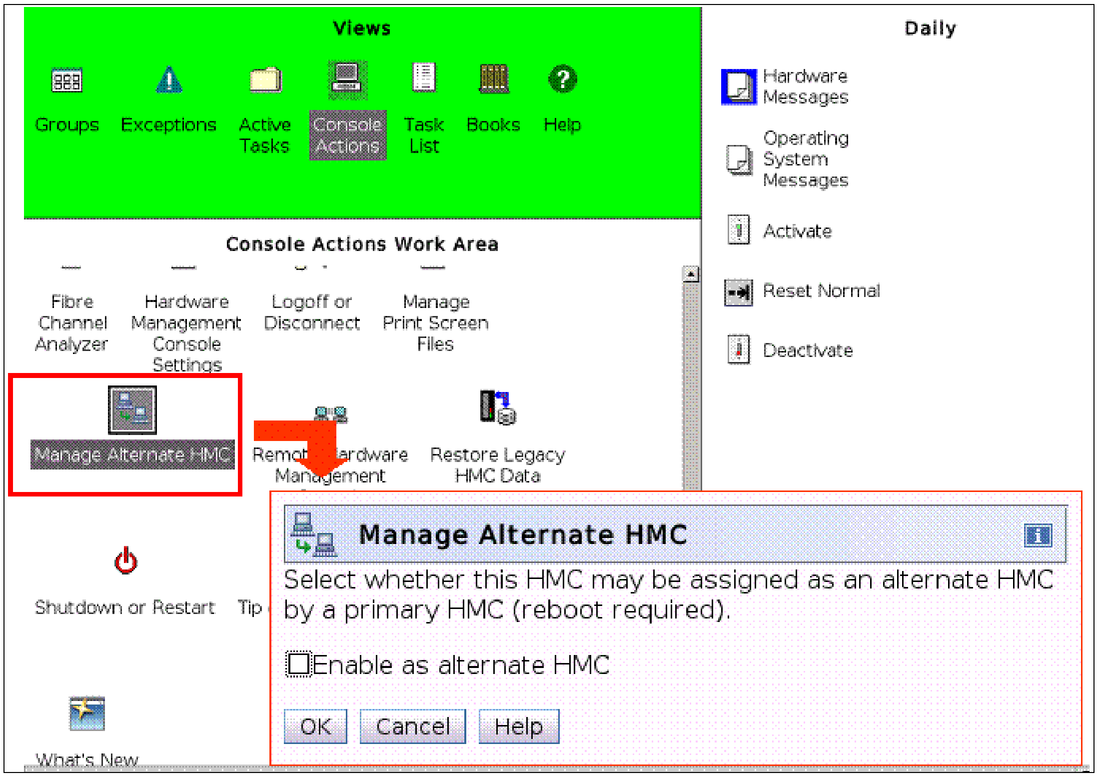

Figure 12-23 Manage Alternate HMC task

Select the alternate HMC from a drop-down list during ensemble creation on the primary HMC. CPCs can then be added as a member of the ensemble during the create ensemble procedure or by using the **Add Member to Ensemble** task. In addition, the **Ensemble Management Guide** task can be used for all tasks that are related to ensemble.

After you create the ensemble, start the initial mirror process between the primary HMC and the alternate HMC.

The following high-level steps are the tasks for setting up an ensemble (for details of these steps, see *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921):

1. Define an alternate HMC.
2. Create the ensemble.
3. Add a member to the ensemble.
4. Install zBX blades.
5. Create and manage storage resources.
6. Create virtual servers, including mounting virtual storage, activating the virtual servers, and opening text consoles.
7. Create and manage virtual networks.

All management actions for the ensemble are conducted from the primary HMC. The alternate HMC has a limited set of tasks. When the primary HMC fails, the alternate HMC takes over all functions for the ensemble.

If a CPC is entered into an ensemble, the **CPC Details** task on the SE and the HMC reflects the ensemble name.

The primary HMC provides the management functions for the ensemble by using Unified Resource Manager (Table 12-5).

Table 12-5 Management functions in the HMC with FC 0025

Functions	Description
Platform Performance Management	Manages Ensemble Image workloads, performance policies, and monitor goals
Virtual Network Management	Manages Ensemble server network connectivity on the Intraensemble Data Network
Virtual Server Management	Manages Ensemble Virtual server workloads, performance, and monitor goals
Hypervisor Management	Manages Ensemble Hypervisor workloads, performance, and monitor goals
z/VM Guest Management[a]	HMC access to the services and APIs needed to manage Ensemble z/VM guests
Entitlement Management	Manages Ensemble Entitlement of different levels of Platform Performance Management
Energy Management	Manages Active Energy Manager to access the Ensemble power/thermal data

a. The z13 does not support z/VM integrated systems management.

12.6.2 Unified Resource Manager

The role of the zEnterprise Unified Resource Manager (zManager) is to ensure that workloads that run on the zEnterprise Node are managed according to specific workload policies. Unified Resource Management manages a single pool of resources, integrating system and workload management across the ensemble environment.

The Unified Resource Manager is a set of functions for system management that can be grouped as follows:

- Defining and managing virtual environment. This function includes automatic discovery and the definition of I/O and other hardware components across IBM z13 and zBX, and the definition and management of LPARs, virtual machines, and virtualized LANs.
- Defining and managing workloads and workload policies.
- Receiving and applying corrections and upgrades to the Licensed Internal Code.
- Running temporary and definitive IBM z13 capacity upgrades.

These functions, which pertain to an ensemble, are provided by the HMC and SEs by using the INMN. All management actions for the ensemble are conducted from the primary HMC.

Because the Unified Resource Manager handles a large amount of definitions and configuration data that are critical to the operability of an ensemble, this information must always be available. This drives the requirement for the primary and alternate ensemble HMC.

Figure 12-24 shows a Unified Resource Manager function in Ensemble. For more information, see *Introduction to Ensembles*, GC27-2609.

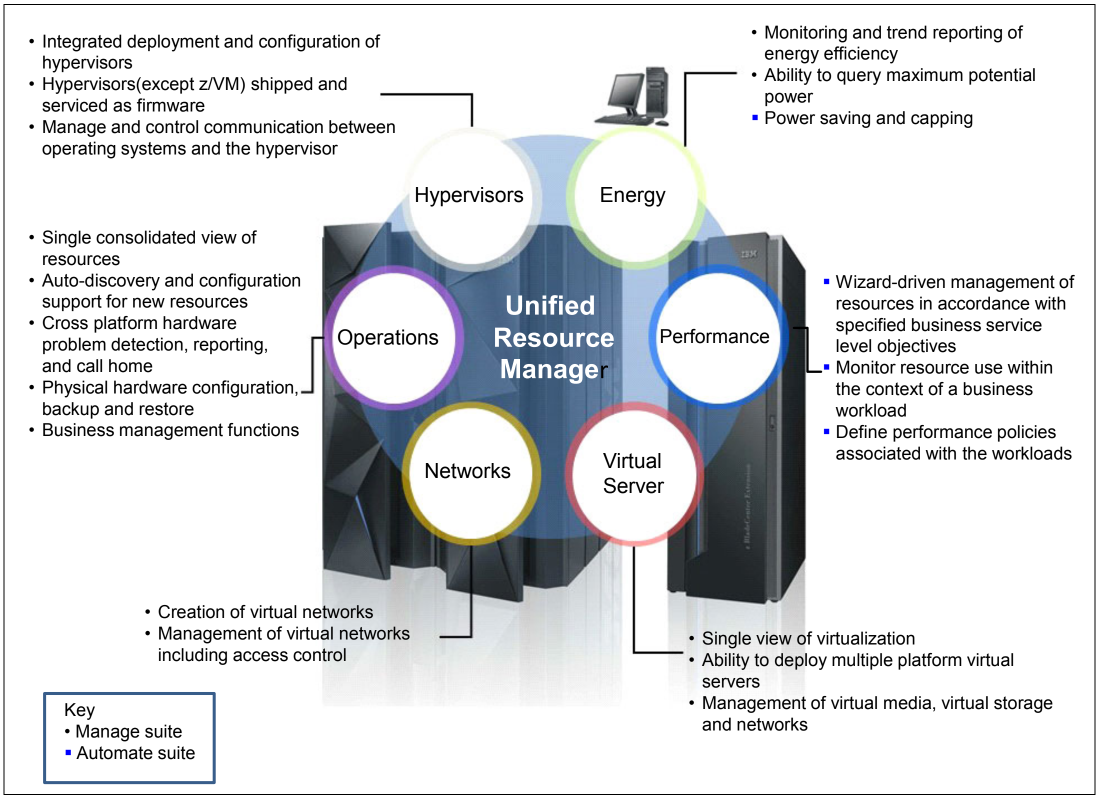

Figure 12-24 Unified Resource Manager functions

The following functions are supported:

- Hypervisor management

  Provides tasks for managing the hypervisor lifecycle, managing storage resources, performing reliability, availability, and serviceability (RAS) and first-failure data capture (FFDC), and monitoring the supported hypervisors.

- Ensemble membership management

  Provides tasks for creating an ensemble and controlling membership of the ensemble.

- Storage management

  Provides a common user interface for allocation and deallocation of physical and virtual storage resources for an ensemble.

- Virtual server management

  Provides lifecycle management to create, delete, activate, deactivate, and modify definitions of virtual servers.

- Virtual network management

  Provides for the management of networking resources for an ensemble.

- Performance management

  Provides a global performance view of all the virtual servers that support workloads that are deployed in an ensemble. The virtual server workload performance goal is like a simplified z/OS WLM policy:

  – You can define, monitor, report, and manage the performance of virtual servers based on workload performance policies.

  – Policies are associated to the workload:

    • From the overall Workload “performance health” report, you can review contributions of individual virtual servers.

    • You can manage resources across virtual servers within a hypervisor instance.

- Energy management

  Monitors energy usage and control power-saving settings, which are accessed through the monitors dashboard. Monitoring of virtual server resources for CPU use and delays, with capability of creating a graphical trend report.

Unified Resource Manager supports various levels of system management. A feature determines the management functions and operational controls that is available for all nodes in an ensemble. These named features are *Manage* and *Automate/Advanced Management*:

- Manage

  Provides Unified Resource Manager’s function for core operational controls, installation, and energy monitoring. It is configured by default and activated when an ensemble is created.

- Automate/Advanced Management

  The Automate function adds goal-oriented resource monitoring management and energy management for CPC components, POWER7 Blade, DataPower XI50, and X Blade. The Advanced Management functionality for IBM System x Blades delivers workload definition, and performance policy monitoring and reporting.

Table 12-6 lists the feature codes that you can order to enable Unified Resource Manager. To get ensemble membership, be sure that you also order FC 0025 for the IBM z13.

Table 12-6 Unified Resource Manager feature codes and charge indicator

Unified Resource Manager managed component	Manage[a] (per connection)	Advanced Management[a] (per connection)	Automate[a] (per connection)
Base features	0019[b] - N/C	N/A	0020[c] - N/C
IFL	N/C	N/A	0052 - Yes
POWER7 Blade	0041 - Yes	N/A	0045 - Yes
DataPower Blade XI50	0040 - Yes	N/A	0044 - N/C
IBM System x Blades	0042 - Yes	0046 - Yes	N/A

a. Yes = charged feature, N/C = no charge, N/A = not applicable. All components are either managed through the Manage suite, or the Automate/Advanced Management suite. The Automate/Advanced Management suite contains the functions of the Managed suite.

b. Feature code 0019 is a prerequisite for feature codes 0020, 0039, 0040, 0041, and 0042.

c. Feature code 0020 is a prerequisite for feature codes 0043, 0044, 0045, 0046, and 0052.

12.7 Additional references

The following IBM publications provide more details about zBX and the Ensemble:

- *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921
- *IBM z13 Technical Introduction*, SG24-8250
- *IBM z13 Technical Guide*, SG24-8251
- *Introduction to Ensembles*, GC27-2609
- *Ensemble Planning Guide*, GC27-2631
- *Ensemble Workload Resource Group Management Guide*, GC27-2633

13

Flash Express

This chapter provides a brief overview and an example of configuring the Flash Express feature.

This chapter includes the following topics:

- Overview of Flash Express
- Installation of Flash Express
- z/OS support for Flash Express
- Linux on z Systems support for Flash Express
- CFCC support for Flash Express
- Manage Flash allocation

For more information about the Flash Express feature, see the following publications:

- *IBM z13 Technical Introduction*, SG24-8250
- *IBM z13 Technical Guide*, SG24-8251

13.1 Overview of Flash Express

Flash Express is internal storage that is implemented with solid-state drives (SSDs) mounted in PCIe Flash Express feature cards. Flash Express feature cards are ordered in pairs for availability. A pair provides 1.4 TB of usable storage. A maximum of four pairs (4 x 1.4 TB = 5.6 TB) are supported in a zEC12 and a z13. Flash Express storage is allocated to each partition similar to main memory allocation. The allocation is specified at the Support Elements (SE) or the Hardware Management Console (HMC).

z/OS uses the Flash Express feature as storage-class memory (SCM) for paging and SVC dumps. SCM is the virtualization of the physical storage across partitions into an abstracted storage space. This allows each logical partition to be configured with its own SCM address space, and allocation by amount (not card size). It also allows you to change the underlying hardware technology while preserving the SCM API. The z/OS paging subsystem supports a mix of Flash Express storage and External Disks.

The Flash Express feature is supported on a z13 server as follows:

- The Flash Express feature can be carried forward with an upgrade from a zEC12 server (FC 0402), and it can be ordered as a new feature for a a z13 server (FC 0403). Both features come in pairs of two cards. The minimum order is two cards (one pair) and the maximum order is eight cards (four pairs).
- Each Flash Express card occupies one I/O slot in a zEC12 or a z13 PCIe I/O drawer. Therefore, one pair occupies 2 I/O slots.
- One PCIe I/O Drawer can support a maximum of two pairs with each Flash Express card installed in a unique domain.
- A second PCIe I/O Drawer is required to support more than two Flash Express pairs.
- Placement of Flash Express cards in PCIe I/O drawers depends on the number of cards that are ordered and the number of PCIe I/O drawers available. Placement preference is for slots 01 and 14 across available PCIe drawers, and then slots 25 and 33.
- Flash Express slot positions are reserved and should be used only when there are no spare slots available.
- The Flash Express feature has no channel-path identifier (CHPID) assigned. Each Flash Express card has a fixed physical channel ID (PCHID) assigned to it defined by its physical location in the drawer.
- No need exists to define a CHPID for the Flash Express feature in the Hardware Configuration Definition (HCD) and I/O configuration program (IOCP). Be careful to avoid the use of Flash Express associated PCHIDs by another device in the HCD IOCP definition.
- Flash Express subchannels are predefined, and are allocated from the 0.25K reserved in subchannel set 0.
- Flash Express is accessed by using the z Systems architected Extended Asynchronous Data Mover (EADM) Facility. Access to Flash Express is initiated with a Start Subchannel instruction.
- Hardware encryption is included in the Flash Express cards with key management using smart cards that are inserted into a smart card reader, which is USB attached to the SE.
- Flash Express cards are connected to each other through *two external cables* (provided by the feature order) to form a RAID 10 Mirror for redundancy.

13.2 Installation of Flash Express

Before Flash Express can be used by z/OS, tasks must be completed by the IBM System Service Representative and the user.

13.2.1 IBM System Service Representative configuration tasks

The IBM System Service Representative performs the following tasks for installation of the Flash Express feature:

- Installs the required MCL bundles.
- Installs smart cards into the smart card readers attached to the primary and alternate Support Elements.
- Installs Flash Express PCIe cards.
- Pairs and formats the Flash Express PCIe cards.

13.2.2 User tasks

After the hardware installation is completed, you are responsible for the following tasks:

- Assigning Flash Express storage to LPARs.
- Configuring the z/OS LPARs to use the assigned Flash Express storage.

These tasks are detailed in the following sections.

13.3 z/OS support for Flash Express

For exploitation of the Flash Express storage z/OS 1.13, an additional Real Storage Management (RSM) Enablement Offering or z/OS 2.1 is required.

There are currently two exploiters under z/OS for Flash Express storage: paging (for local data sets) and SVC dumps. To use Flash Express for paging, you must configure Flash Express storage online to z/OS. SVC dump processing automatically uses Flash Express storage after it is online to z/OS. No extra commands or setup are required for SVC dumps.

Consider the following information when placing data on SCM storage and page data sets:

- Pageable Link Pack Area (PLPA): At IPL time, PLPA pages are placed on both SCM storage and disk.
- Virtual input/output (VIO): VIO data is always placed on disk. Local page data sets are required for VIO, and when needed to support peak paging demands that require more capacity then provided by the amount of configured SCM storage.
- Pageable large pages: If contiguous SCM space is available, pageable large pages are written to SCM storage. If SCM storage is not available, pageable large pages are backed with 4 KB page frames.
- All other data: If available space exists on both SCM storage and disk, the selection (where to place the data) is based on the response time of SCM storage rather than disk.

Reminder: You should always review the PSP Device buckets for the recent z/OS requirements. Any requirements stated in this document might not be as current.

13.3.1 z/OS V1R13 RSM Enablement Offering web deliverable

For z/OS 2.1, exploitation of Flash Express storage is included in the base support.

For z/OS 1.13, the required enhancements for the RSM for using Flash Express are delivered in the z/OS V1R13 RSM Enablement Offering web deliverable (FMID JBB778H and JBB778K). With RSM Enablement Offering and Flash Express, z/OS is also designed to help improve processor performance by supporting middleware such as IMS, with its exploitation of pageable large (1 MB) pages.

Additional follow-on support is now available with APARs OA40967 and OA40968:

- OA40967 provides support for 2 GB pages through enhancements to the IARV64 service.
- OA40968 provides support to configure SCM increments offline through a z/OS operator command, and allows the PLPA and COMMON paging data sets to be optional.

Exploitation of 1 MB pages is available as follows:

- z/OS V1.13 IBM Language Environment® when used with a run-time option.
- Java, with the IBM 31-bit SDK for z/OS, Java technology Edition, V7.0.0 (5655-W43) and SDK7 IBM 64-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W44).
- The IMS Common Queue Server, which is designed to use pageable large pages for selected buffers when running IMS 12 (5635-A03) on an zEC12 server with the PTF for APAR PM66866.

Exploitation of 2 GB pages is available as follows:

- IBM 31-bit SDK for z/OS, Java technology Edition, V7.0.0 (5655-W43) and SDK7 IBM 64-bit SDK for z/OS, Java Technology Edition, V7.0.0 (5655-W44).

IBM RMF™ uses the term storage-class memory (SCM) as a synonym for Flash Express memory. With RMF function APAR OA38660, RMF provides the SCM and pagegable large pages support for z/OS 1.13. The support enhances RMF Postprocessor, Monitor II, and Monitor III reports with various new statistics for SCM and pageable large pages.

New SCM statistics are provided in these reports:

- RMF Postprocessor Paging Activity report
- RMF Postprocessor Page Data Set Activity (PAGESP) report
- RMF Monitor II Page Data Set Activity (PGSP) report

New statistics for pageable large pages are provided in these reports:

- RMF Postprocessor Paging Activity report
- RMF Postprocessor Virtual Storage Activity (VSTOR) report
- RMF Monitor III Storage Memory Objects (STORM) report

A program directory for the RSM Enablement Offering for z/OS V1R13 is available:

`http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=ZSL03261USEN`

13.3.2 z/OS system parameter and command enhancements

An IEASYSxx parameter, PAGESCM, is added for z/OS V1R13. Additionally, several display commands and the **config** command are enhanced on z/OS for Flash Express support. Becoming familiar with these facilities before you use Flash Express is important.

For Flash Express, the PAGESCM parameter is supported in IEASYSxx. The syntax is shown in Example 13-1. This parameter determines whether and how much Flash Express storage is made available to an LPAR at IPL time.

Example 13-1 PAGESCM parameter

```
PAGESCM={xxxxxxM              }
        {xxxxxxG              }
        {xxT                  }
        {ALL                  }
        {NONE                 }
        {0                    }
```

This parameter specifies the minimum amount of storage-class memory (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. Additional SCM will be allocated on an as-needed basis if use of this initial amount of SCM is exceeded.

You may 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 will be reserved for paging at IPL. Instead, SCM is allocated as needed, based on paging demand.

Default value ALL

Associated parmlib member: None

Additionally an enhancement is made to the **`CONFIG`** command. The **`CONFIG SCM`** command is used to configure storage-class memory (SCM) storage online or offline to an LPAR (Example 13-2).

Example 13-2 CONFIG SCM

```
CONFIG SCM(ddddddddM|G|T),ONLINE|ON
CONFIG SCM(ddddddddM|G|T),OFFLINE|OFF
CONFIG SCMscm_ranges,OFFLINE|OFF
```

The system reconfigures SCM, both logically and physically. To bring SCM online, an amount must be specified. To take SCM offline, a range of starting and ending addresses of the SCM blocks must be specified.

The command has the following values:

ddddddddM|G|T The amount of storage-class memory (SCM) to be reconfigured. Specify up to eight decimal digits followed by a multiplier (M=megabytes, G=gigabytes, T=terabytes) for this amount. Check your processor configuration for the supported SCM increment sizes. The value for `dddddddd` must be a multiple of the SCM increment size (usually 2, 4, or 8), and cannot exceed `16T`.

Instead of specifying a decimal amount, you can alternatively specify a hexadecimal amount, with or without a multiplier, in the following format:

X'xxxxxx'

For example:

X'123456789A00000'

X'123'M

You can use underscore characters in any hexadecimal specification for additional 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 will select which SCM increments to take offline.

ONLINE or ON The system brings the specified amount of SCM online. ONLINE is the default value if only `CONFIG SCM` is specified. The system rejects the command if you specify these values:

– A value that is not a multiple of the SCM increment size.

– A value that exceeds the total amount of SCM that is defined to this partition.

– A value that is not a valid amount of SCM (0, for example).

SCM is not supported or not defined on the system.

OFFLINE or OFF The system takes the specified amount or specified ranges of SCM offline.

> **Attention:** Taking SCM offline can affect data reliability and performance. Consider these implications before taking SCM offline:
>
> - Your system must have enough auxiliary storage, which can include SCM and must include page data sets, to back critical system data. The **`CONFIG SCM OFFLINE`** command will fail 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 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 there is an insufficient amount of SCM, 1 MB large pages are demoted to 256 4-KB pages and paged 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 `D ASM` and `D M` commands are enhanced to display information and status that are related to Flash Express:

`D ASM`	Lists SCM status along with paging data set status.
`D ASM,SCM`	Displays a summary of SCM usage.
`D M=SCM`	Displays SCM online/offline and increment information.
`D M=SCM(DETAIL)`	Displays detailed increment-level information.

Tip: You might notice a difference in usage numbers between the `D M=SCM` and `D ASM` commands. The difference is because of how ASM perceives its use of the cache of available SCM blockids that ASM maintains. To ASM, some blockids are not in use because they were not yet assigned to page-out requests. However, to the `D M=SCM` command processor, blockids are in use because they were assigned to ASM for its use.

13.4 Linux on z Systems support for Flash Express

Flash Express exploitation on zEC12 and zEC13 is also available for the following Linux on z Systems distributions:

- SLES 11 SP3 or later
- RHEL 6.4 or later

13.5 CFCC support for Flash Express

With z/OS 1.13 with the RSM Enablement Offering or with z/OS 2.1 running in every connected z/OS LPAR, Flash Express storage can also be exploited by coupling facility LPARs running CFCC Level 19 (on zEC12) or CFCC Level 20 (on zEC13). Systems without this support cannot connect to or rebuild a structure using Flash Express storage.

The allocation of Flash Express storage to a coupling facility LPAR is done in the same way as for z/OS LPARs and is described in 13.6, “Manage Flash allocation” on page 638. The amount of SCM allocated to a coupling facility LPAR can be displayed in the Operating System Messages panel at the HMC. For example, LPAR A0E, which has allocated 16 GB of flash memory (Figure 13-2 on page 639), has a message that shows the amount of SCM available, as shown in Figure 13-1.

```
CF0280I CFCC Release 20.00, Service Level 00.10
        Built on 09/09/2014 at 10:33:00
        Code Load Features:
           Facility Operational Level: 20
CF0011I Coupling Facility is active with:
           2 CPs
           4 CF Receiver Channels
           4 CF Sender Channels
           7716 MB of allocatable storage
           16384 MB of Total SCM storage
```

Figure 13-1 CFCC messages with SCM

The CF must know the algorithm of how the structure will be used by the application. Currently, this is defined only for IBM MQ shared queues. To use this function, assign flash memory to your coupling facilities according to the procedure described in the following section and update your structure definitions in your CFRM policy with the new parameter SCMMAXSIZE and SCMALGORITHM. For details see *z/OS MVS Setting Up a Sysplex*, SA23-1399.

IBM MQ for z/OS Version 7 (5655-R36) allows the migration of IBM MQ shared queue objects to flash memory when structure utilization exceeds the defined threshold. The IBM MQ objects are fetched back to real CF Storage when requested. This provides an overflow capability for IBM MQ shared queues to handle workload peaks.

With RMF new function APAR OA40515, RMF provides measurement data and reporting capabilities for Flash Express. The support enhances RMF Postprocessor and Monitor III reports with various new CF SCM statistics.

New Coupling Facility SCM statistics are provided in these reports:

- RMF Postprocessor Coupling Facility Activity (CF) report
- RMF Monitor III Coupling Facility Overview (CFOVER) report
- RMF Monitor III Coupling Facility Activity (CFACT) report

13.6 Manage Flash allocation

Assignment of Flash Express storage to LPARs is done through the HMC or SE by using the Manage Flash Allocation window (Figure 13-2 on page 639), from the Configuration task. In this window, you can allocate Flash Express storage in increments of 16 GB to any partition defined in these I/O configuration data sets (IOCDS): A0, A1, A2, and A3. z/OS currently supports up to 16 TB of flash memory in a single system image.

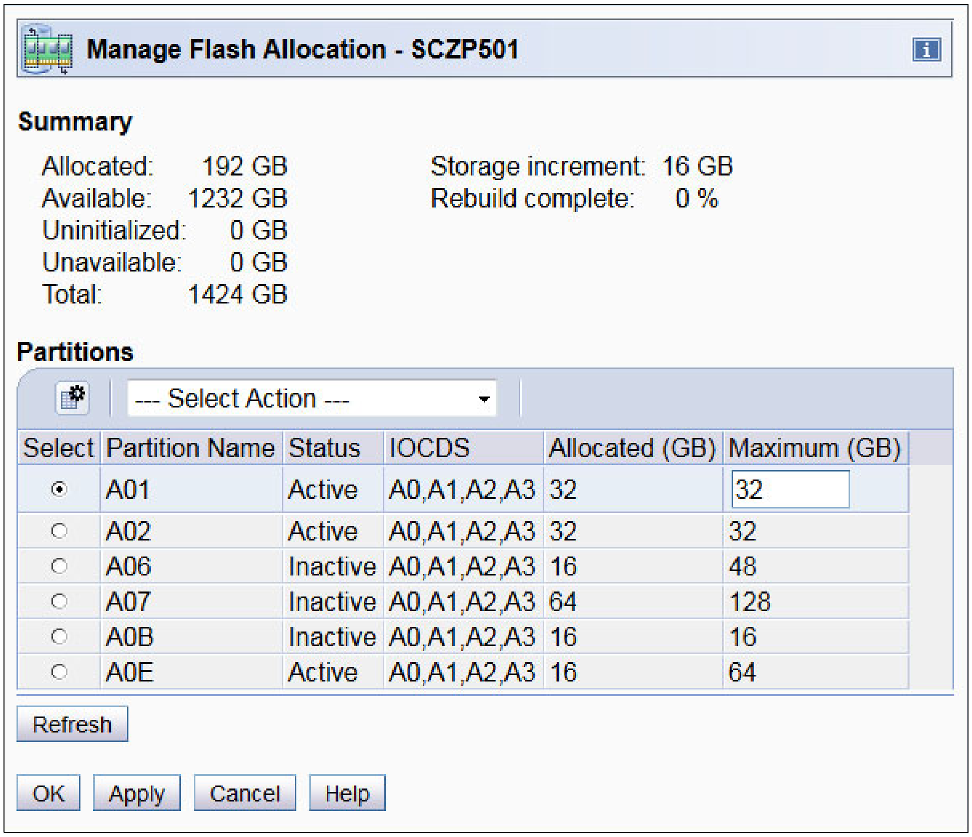

Figure 13-2 Manage Flash Allocation window

A summary view is initially presented. It displays information about the Flash Express storage allocation. The summary includes these fields:

Allocated Amount of Flash Express allocation that is currently assigned to the CPC

Available Amount of Flash Express allocation that is currently available on the CPC

Uninitialized Amount of Flash Express allocation that is currently uninitialized on the CPC

Unavailable Amount of Flash Express allocation that is currently unavailable on the CPC

Total Current total amount of all Flash Express allocations on the CPC

Storage increment Current maximum Flash Express storage increment on the CPC

Rebuild complete Completion percentage of the new Flash Express allocations rebuild

The partition table lists partitions and current Flash Express storage allocation information. The table includes the following fields:

Partition Name Name of the logical partitions that are assigned to the CPC with Flash Express allocations.

Status Status of the logical partitions. The status can be *Active* or *Inactive*.

IOCDS Logical partition names that have Flash Express allocations.

Allocated (GB) Current amount of Flash Express storage that is allocated for the logical partitions. Select the logical partition to change the initial Flash allocation increment.

Maximum (GB) Maximum amount of Flash Express storage that is allocated for the logical partitions. Select the logical partition to change the maximum Flash allocation increment.

Use the **Select Action** menu to modify allocations to selected partitions or to view the logical partition to PCHID assignments. Figure 13-3 shows an example of the view partition to PCHID map display.

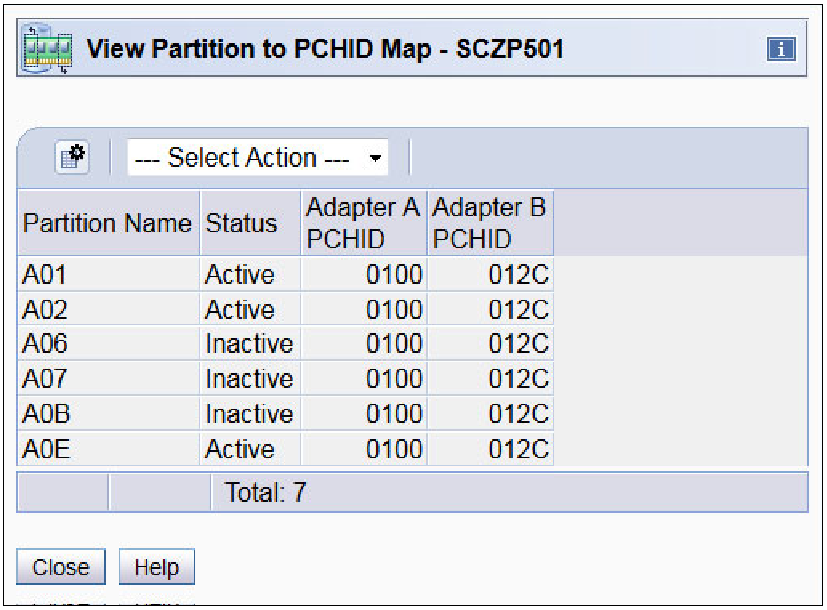

Figure 13-3 View Partition to PCHID Map

13.6.1 Allocating Flash Express storage to partitions

Be aware of these considerations when you allocate Flash Express storage to a partition:

- When an allocation is first defined, you must set the initial and maximum allocation in 16 GB increments.
- A new allocation will be detected by an active z/OS LPAR, but no SCM memory will be varied online to z/OS.
- An 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 defined in this window.
- Any extra memory varied online or offline by z/OS is reflected in the allocated GB in this window if refreshed.
- Minimum amounts are allocated from the available pool, so they cannot be overallocated.
- Maximum amounts can be overallocated up to 5696 GB.
- Maximum amounts must be greater than or equal to the initial amounts.

To allocate Flash Express storage to a partition, click **Select Action** → **Add allocation**. The New Flash Allocation panel opens (Figure 13-4).

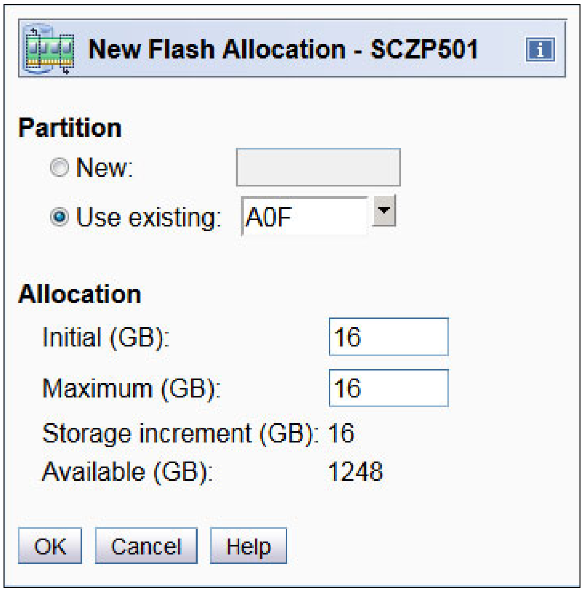

Figure 13-4 New Flash Allocation

Enter the partition manually by selecting **New** and providing the partition name; or choose a partition by selecting **Use existing** and then selecting the partition from the drop-down menu. Only partitions that do not have an existing allocation are shown in this drop-down menu. The new allocation window cannot be used to modify an existing allocation for a partition.

Enter allocation values for the specified partition name in the following fields:

Initial (GB) Enter the initial Flash allocation to be used for the logical partition.

Maximum (GB) Enter the maximum Flash allocation to be used for the logical partition.

The following fields display information:

Storage Increment (GB) Displays the Flash increment value.

Available (GB) Displays the amount of flash memory that is currently available.

13.6.2 Allocation example

The example in this section allocates Flash Express storage for the active z/OS image named SC76 in the partition named A03 (Channel Subsystem 0, Multiple Image Facility ID 3). The RSM Enablement Offering support is installed.

From SC76, issue the MVS **`D IPLINFO,PAGESCM`** command. Example 13-3 shows the results.

Example 13-3 D IPLINFO,PAGESCM

```
D IPLINFO,PAGESCM
IEE255I SYSTEM PARAMETER 'PAGESCM': NOT_SPECIFIED
```

Because no `PAGESCM` parameter is specified, the default value of ALL is used (reserves all SCM for paging at IPL time). If a Flash Express storage allocation is defined for the LPAR and

`PAGESCM=ALL` is specified (or kept at the default), then at IPL time, the initial amount of Flash Express storage specified is used automatically by z/OS for paging. Likewise, if a specific amount is specified, this amount is made available for paging.

However, if a Flash Express allocation is added when a z/OS image is active, z/OS detects the allocation, but does not automatically vary the storage online for use. The **`CONFIG SCM`** MVS operator command must be issued to vary the Flash Express storage online for use as paging space for the z/OS LPAR.

When you dynamically allocate and configure SCM storage online to z/OS, z/OS uses SCM for local paging. You will still see percentages for the local page data sets. After IPL, z/OS uses SCM primarily for local paging, so your local page data sets should be 0% or a lot lower than SCM storage.

To add a Flash Storage allocation for partition A03, from the Manage Flash Allocation task, select an existing partition, and click **Select Action** → **Add allocation**. The New Flash Allocation window opens. Click **Use existing** and select partition **A03**. Enter the Initial and Maximum amounts of 16 GB, as shown in Figure 13-5. Then, click **OK** to define the allocation.

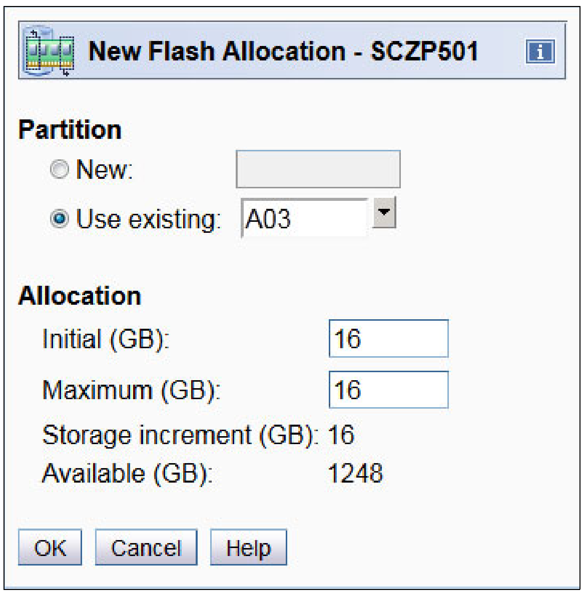

Figure 13-5 Flash Allocation example for partition A03

A window opens to indicate that the allocation was successfully added (Figure 13-6). Click **OK**.

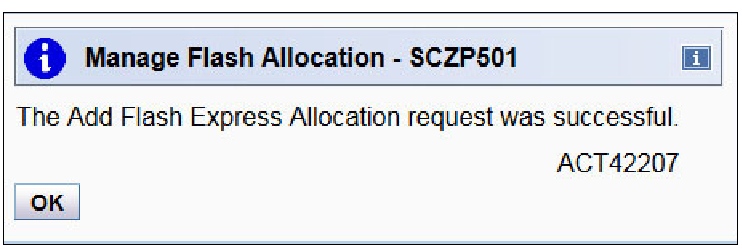

Figure 13-6 Flash Express allocation confirmation

The new allocation is displayed on the Manage Flash Allocation Partitions table (Figure 13-7).

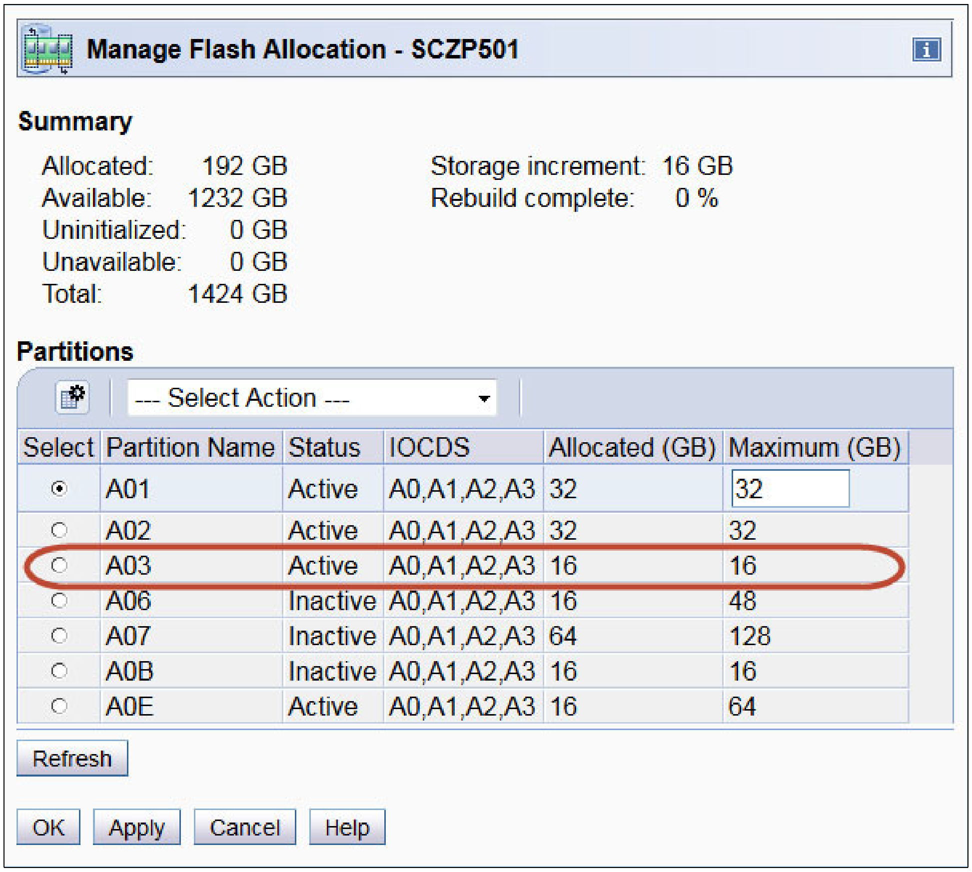

Figure 13-7 Updated partitions table

When the allocation is defined through the HMC or SE for an active z/OS image, the allocation is detected. On the z/OS image's (SC76) console, the IAR034I message is displayed:

```
IAR034I ADDITIONAL STORAGE-CLASS MEMORY DETECTED
```

From SC76, run the enhanced **D ASM** and **D M** commands to display Flash Express SCM-related information and status. The result for each command is shown in Example 13-4.

Example 13-4 Display commands

```
D ASM
IEE200I 13.53.38 DISPLAY ASM 976
TYPE     FULL STAT   DEV  DATASET NAME
PLPA     100% FULL  1520  PAGE.SC76.PLPA
COMMON    99%   OK  1520  PAGE.SC76.COMMON
LOCAL      0%   OK  1520  PAGE.SC76.LOCAL1
LOCAL      0%   OK  DE0A  PAGE.SC76.LOCAL2
LOCAL      0%   OK  DF09  PAGE.SC76.LOCAL3
PAGEDEL COMMAND IS NOT ACTIVE

D ASM,SCM
IEE207I 13.54.14 DISPLAY ASM 978
STATUS      FULL                SIZE             USED          IN-ERROR
NOT-USED
```

```
D M=SCM
IEE174I 13.55.04 DISPLAY M 980
STORAGE-CLASS MEMORY STATUS
16G DEFINED
16G OFFLINE-AVAILABLE
SCM INCREMENT SIZE IS 16G

D M=SCM(DETAIL)
IEE174I 13.56.26 DISPLAY M 982
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
16G DEFINED
ONLINE: 0G  OFFLINE-AVAILABLE: 16G  PENDING OFFLINE: 0G
SCM INCREMENT SIZE IS 16G
```

From these commands, you see that 16 GB of Flash Express storage is available (defined), but not in use (offline-available).

To vary the storage online to the example LPAR, issue the **CONFIG SCM(xxG),ONLINE** command, as shown in Example 13-5 along with results. The amount of storage configured online must be specified according to the supported increment size. From these displays, the supported increment size is 16G. Because you have only allocated 16G for this z/OS image, vary the entire (16G) amount online.

Example 13-5 CF SCM(16G),ONLINE

```
CF SCM(16G),ONLINE
IEE195I SCM LOCATIONS 0G TO 16G ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Issue the **D ASM** and **D M** commands again to display the status of the Flash Express storage and see that the 16 GB initial value is now online and available (Example 13-6).

Example 13-6 Post configuration displays

```
D ASM
IEE200I 13.59.32 DISPLAY ASM 987
TYPE     FULL STAT   DEV  DATASET NAME
PLPA     100% FULL  1520  PAGE.SC76.PLPA
COMMON    99%   OK  1520  PAGE.SC76.COMMON
LOCAL      0%   OK  1520  PAGE.SC76.LOCAL1
LOCAL      0%   OK  DE0A  PAGE.SC76.LOCAL2
LOCAL      0%   OK  DF09  PAGE.SC76.LOCAL3
SCM        0%   OK   N/A   N/A
PAGEDEL COMMAND IS NOT ACTIVE

D ASM,SCM
IEE207I 14.01.46 DISPLAY ASM 989
STATUS      FULL              SIZE              USED         IN-ERROR
IN-USE        0%         4,194,304                 0                0

D M=SCM
IEE174I 14.02.58 DISPLAY M 994
STORAGE-CLASS MEMORY STATUS
16G DEFINED
ONLINE
```

```
 0G-16G
0G OFFLINE-AVAILABLE
0% IN USE
SCM INCREMENT SIZE IS 16G

D M=SCM(DETAIL)

IEE174I 14.04.33 DISPLAY M 996
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
16G DEFINED
ADDRESS  IN USE  STATUS
     0G      0%  ONLINE
ONLINE: 16G  OFFLINE-AVAILABLE: 0G  PENDING OFFLINE: 0G
0% IN USE
SCM INCREMENT SIZE IS 16G
```

Before adding the Flash Express allocation to this z/OS image (SC76) at IPL time, the message in Example 13-7 is issued. It indicates that although PAGESCM=ALL was specified (kept the default in the example), no SCM memory (Flash Express storage) is available and brought online.

Example 13-7 Preallocation IPL message

```
IAR031I USE OF STORAGE-CLASS MEMORY FOR PAGING IS ENABLED - PAGESCM=ALL
       , ONLINE=00000000M
```

After adding the Flash Express allocation to this z/OS image (SC76), at the next IPL, the next message is issued (Example 13-8).

Example 13-8 Post-allocation IPL message

```
IAR031I USE OF STORAGE-CLASS MEMORY FOR PAGING IS ENABLED - PAGESCM=ALL
       , ONLINE=00016384M
```

The output indicates that PAGESCM=ALL was specified (kept the default in the example) and all available (initial amount specified for the allocation) storage is brought online.

13.6.3 Increasing the available Flash Express storage allocation

After an allocation is defined for a partition, the initial value cannot be changed unless you delete the allocation and define a new allocation for the partition with a different initial value. However, the maximum allocation can be changed.

The maximum allocation for a partition can be changed through the Manage Flash Allocation window. When a partition is selected, the Maximum field can be changed as shown in Figure 13-8. The maximum allocation cannot be less than the initial allocation value.

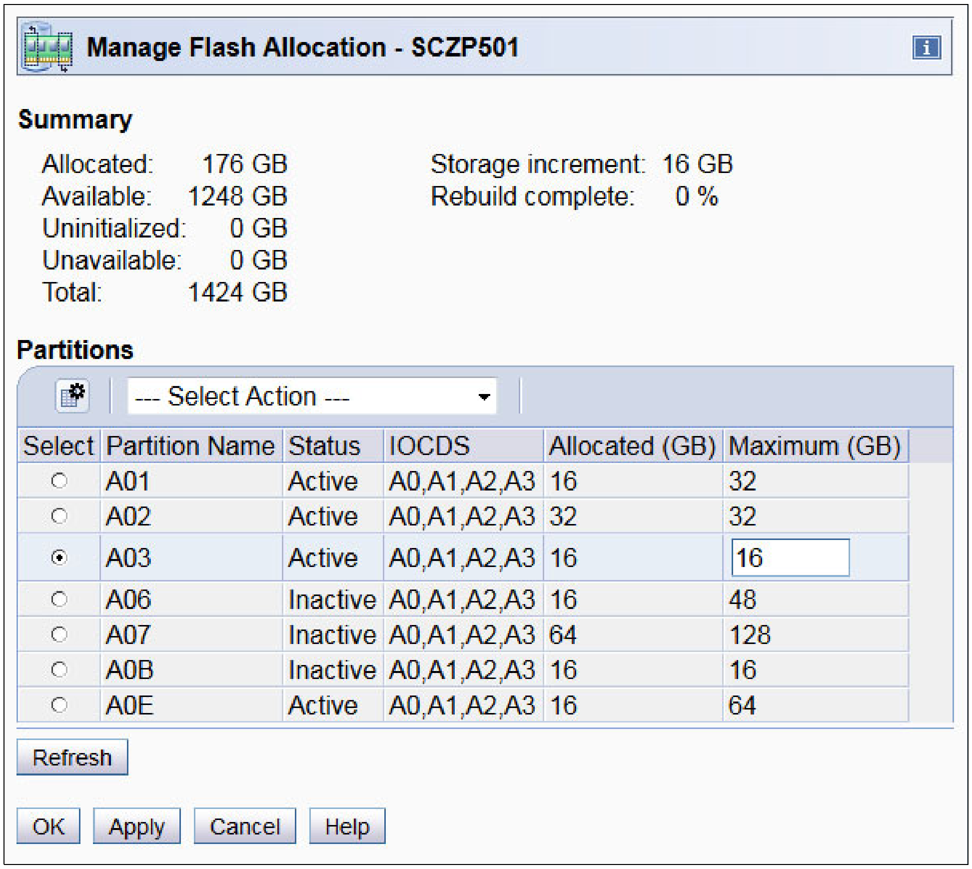

Figure 13-8 Maximum Field overtypable

As an example, increase the maximum Flash Express storage for partition A03 (z/OS image SC76) from 16 GB to 32 GB. Enter `32` in the **Maximum** field as shown in Figure 13-9 on page 647. Click **Apply** to update the allocation definition.

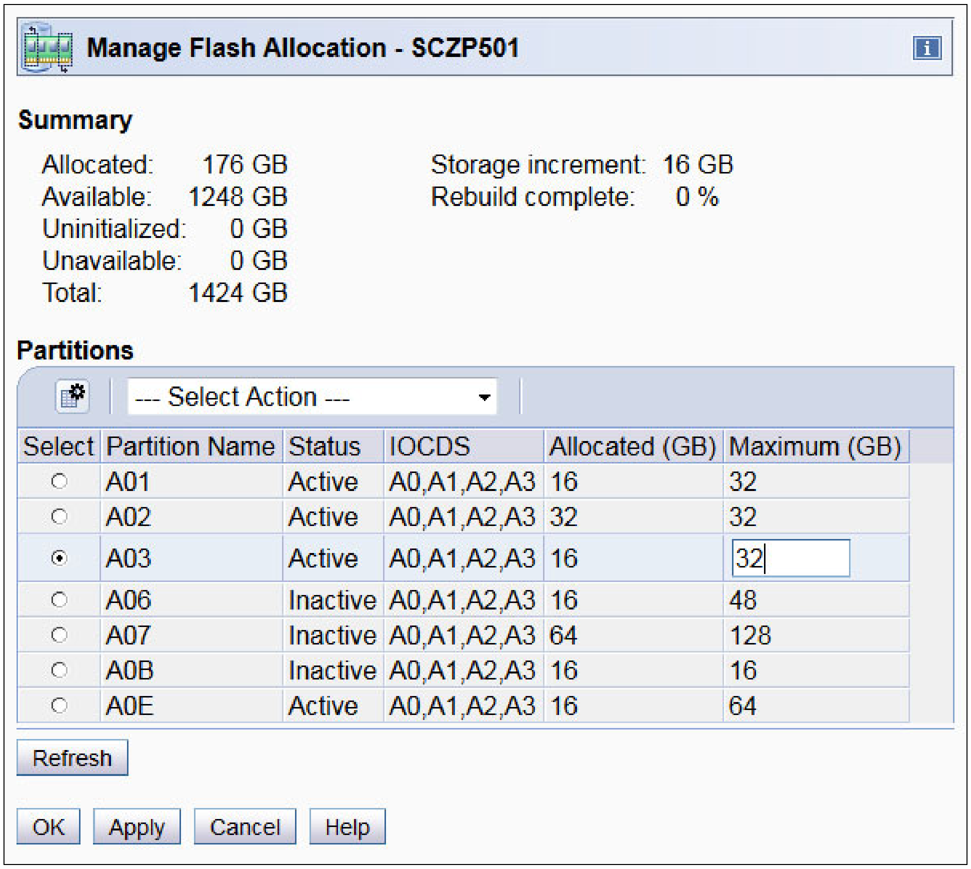

Figure 13-9 Increasing the maximum allocation

A window indicates that the update was successful (Figure 13-10).

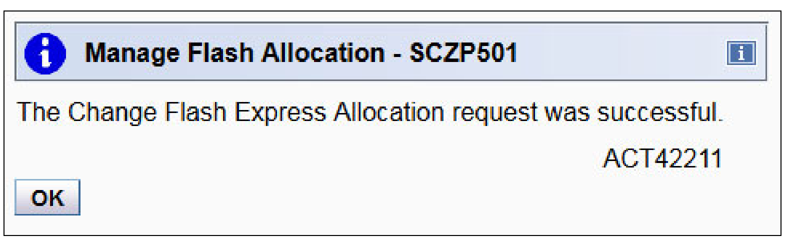

Figure 13-10 Maximum files update successful

The new value is now reflected in the Manage Flash Allocation partition table (Figure 13-11).

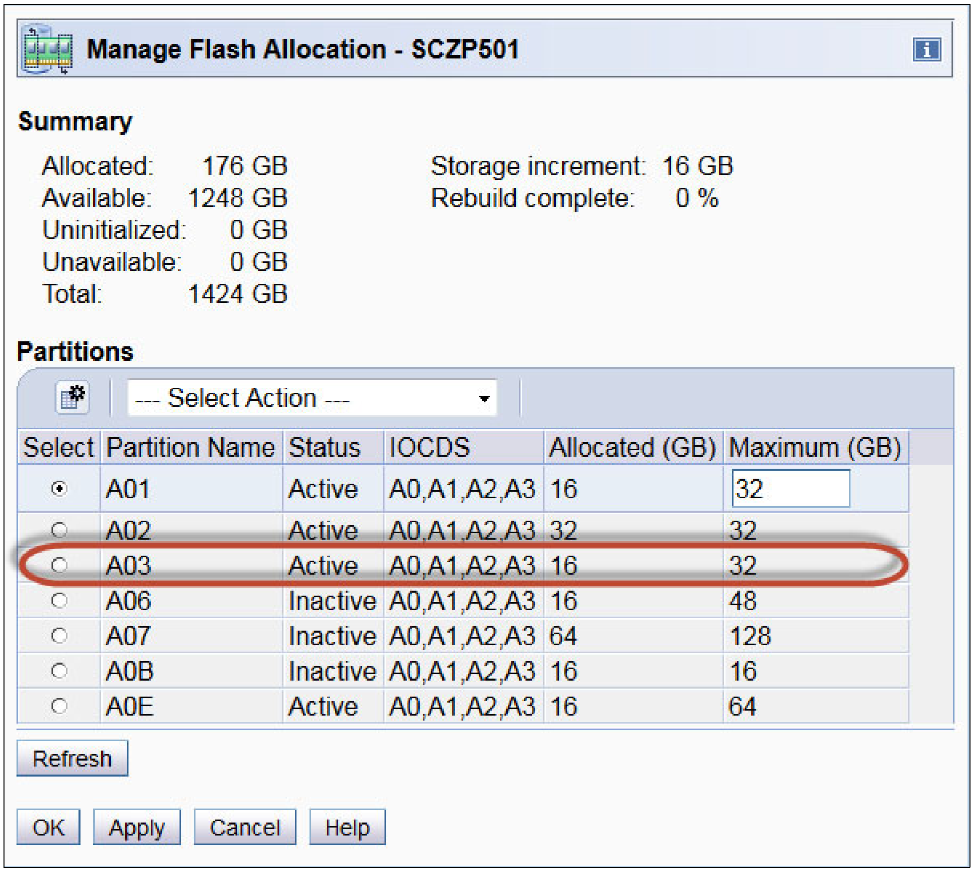

Figure 13-11 Maximum amount updated in partition table

On an active z/OS image, the allocation change will be detected and an IAR0341I message is generated:

```
IAR034I ADDITIONAL STORAGE-CLASS MEMORY DETECTED
```

Issuing a **`D M=SCM(DETAIL)`** command shows that 16 GB of extra SCM storage is available but offline (Example 13-9).

Example 13-9 D M=SCM(DETAIL)

```
D M=SCM(DETAIL)
IEE174I 15.17.19 DISPLAY M 228
STORAGE-CLASS MEMORY STATUS - INCREMENT DETAIL
32G DEFINED
ADDRESS   IN USE   STATUS
     0G       0%   ONLINE
ONLINE: 16G  OFFLINE-AVAILABLE: 16G  PENDING OFFLINE: 0G
0% IN USE
SCM INCREMENT SIZE IS 16G
```

This additional storage can now be configured online by using the **CONFIG SCM(xxG),ONLINE** command (Example 13-10).

Example 13-10 Configure extra SCM storage online

```
CONFIG SCM(16G),ONLINE
IEE195I SCM LOCATIONS 16G TO 32G ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

13.6.4 Removing an allocation

Flash Express storage allocations can be removed from selected partitions through the HMC or SE Manage Flash Allocation window. An allocation cannot be deleted if the partition is active. You can delete and redefine an allocation for a partition if you want to change the initial value.

However, before permanently removing an allocation to a z/OS image that was configured to use SCM storage, consider the following implications:

- Your system must have enough auxiliary storage, in the form of page data sets, to back critical system data.
- SCM is required for paging large pages.

In the following example, an allocation that was defined for partition A0B is removed.

From the Manage Flash Allocation window, select the partition whose allocation is to be deleted, and click **Select Actions** → **Remove allocation** (Figure 13-12 on page 650). The status of the partition must be inactive or the removal will fail. Click **OK**.

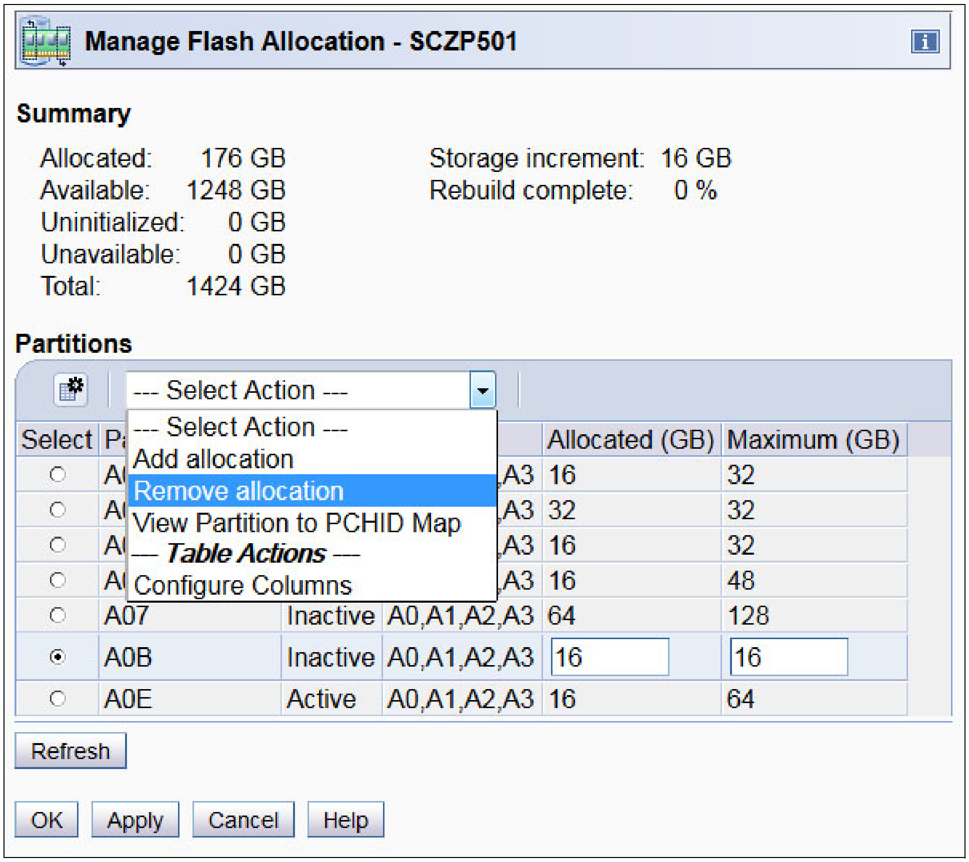

Figure 13-12 Remove allocation

A confirmation prompt and warning message is displayed (Figure 13-13). Click **Yes** to continue with the removal.

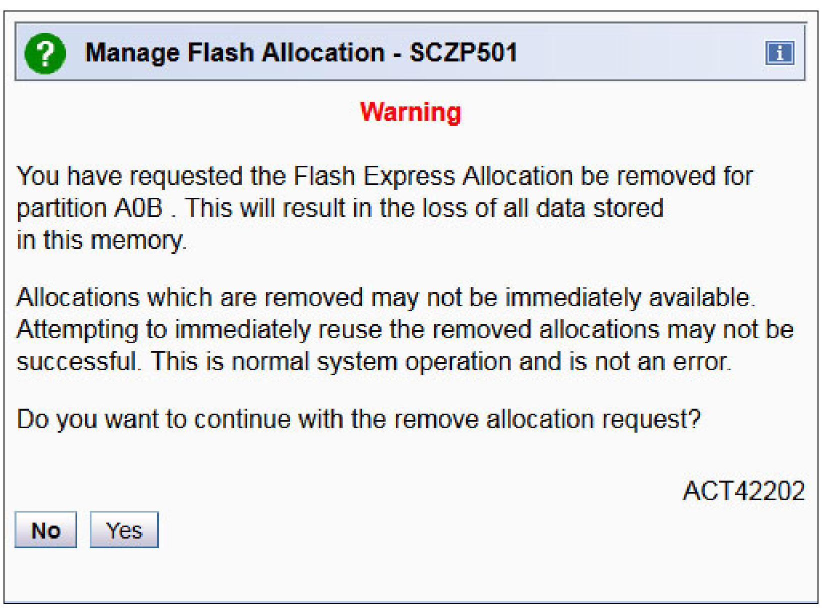

Figure 13-13 Remove allocation warning

A message is displayed, indicating the successful removal of the allocation (Figure 13-14).

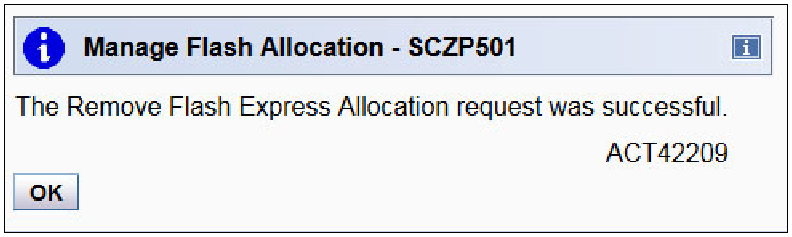

Figure 13-14 Remove allocation success message

The allocation for the partition no longer is displayed in the Manage Flash Allocation partition table. A new allocation for the partition can now be created if you want.

14

Configure Open System Adaptor (OSA) and RoCE Express

This chapter explains how to configure the OSA-Express and RDMA over Converged Ethernet (RoCE) Express features.

This chapter includes the following topics:

- Configure OSE channel with Open System Facility (OSA/SF) on the HMC
- 10GbE RoCE Express HCD definitions and z/OS display commands
- HCD definitions and z/OS display commands for zBX
- z/OS display commands
- Configure OSA-Express Integrated Console Controller (OSA-ICC)
- Setting OSA parameters by using OSA Advanced Facilities

14.1 Configure OSE channel with Open System Facility (OSA/SF) on the HMC

If you define OSA as OSE channel (non-QDIO mode), you must customize OSA with OSA/SF *except* for the following uses.

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

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

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 through the HMC. The OSA CHPID type and OSA/SF requirements are listed in Table 14-1.

Table 14-1 OSA/SF and OSA CHPID reference

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

From zEC12 GA2, OSA/SF is available on HMC. You must use OSA/SF on HMC to define OSA Address Table (OAT) and SNA definition for OSA-Express5S. You can still use OSA/SF running on z/OS for OSA-Express4S, but we suggest that you migrate to OSA/SF on HMC. Table 14-2 is a comparison of OSA/SF.

Table 14-2 OSA/SF comparison

Type of OSA/SF	Supported OSA Express
OSA/SF on HMC	OSA Express-5S, OSA Express-4S
OSA/SF on z/OS	OSA Express-4S

If you are upgrading from z196 or zEC12 to z13, include OSA Express-4S 1000Base-T feature, your OSE configuration will be automatically migrated. We suggest that you check the configuration in OSA/SF on the HMC.

14.1.1 Customize OAT and SNA timer with OSA/SF on HMC

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

Complete these steps:

1. Start OSA/SF on the HMC by opening OSA Advanced Facilities. Select the CPC for the OSA in which you want to customize, then select the **Operational Customization** → **OSA Advanced Facilities** task (Figure 14-1).

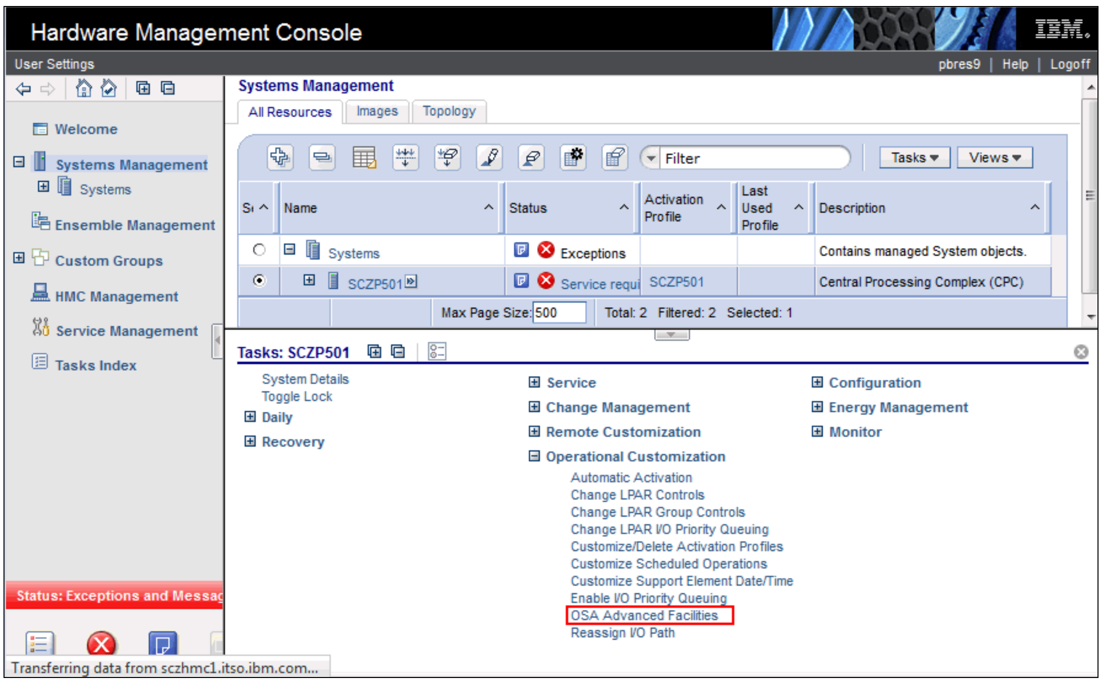

Figure 14-1 OSA/SF on HMC: Selecting OSA Advanced Facilities

2. The OSA Advanced Facilities panel opens (Figure 14-2). Select the PCHID you want to customize. In this example, we selected PCHID 0144. Click **OK**.

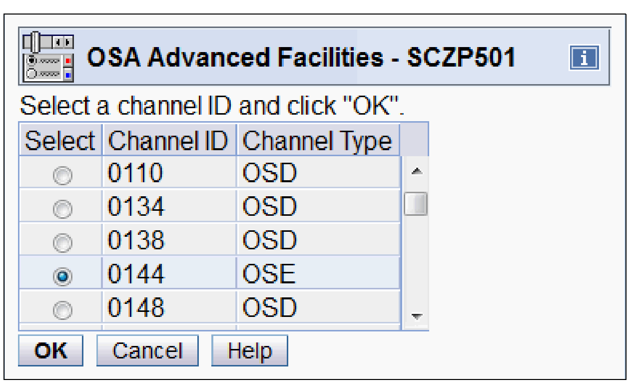

Figure 14-2 OSA/SF on HMC: Selecting CHPID

3. Select **Card specific advanced facilities** and click **OK** (Figure 14-3).

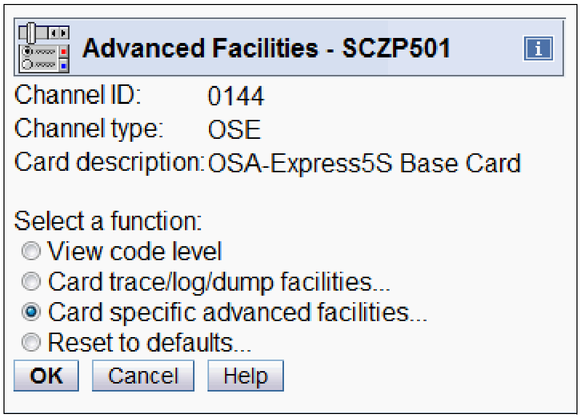

Figure 14-3 OSA/SF on HMC: Card specific advanced facilities selection

4. The Advanced Facilities panel opens(Figure 14-4). OSA/SF on HMC is integrated in this panel. If you edit OAT and SNA timer entries, select **Panel configuration options** and click **OK**.

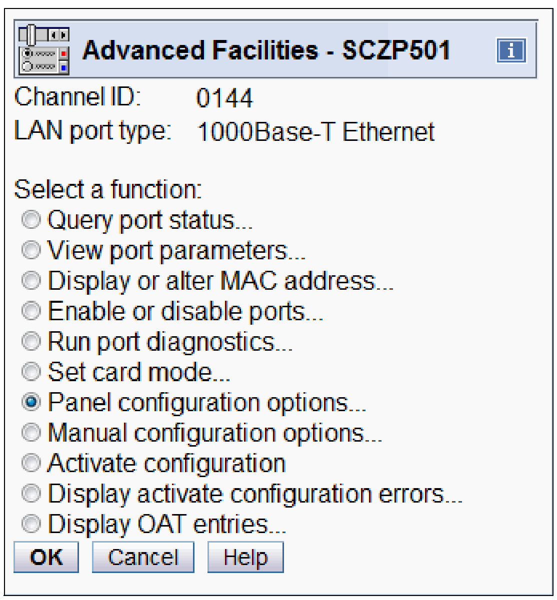

Figure 14-4 OSA/SF on HMC: Panel configuration options selection

5. Panel Configuration Options opens (Figure 14-5 on page 657). You can define these items:

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

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

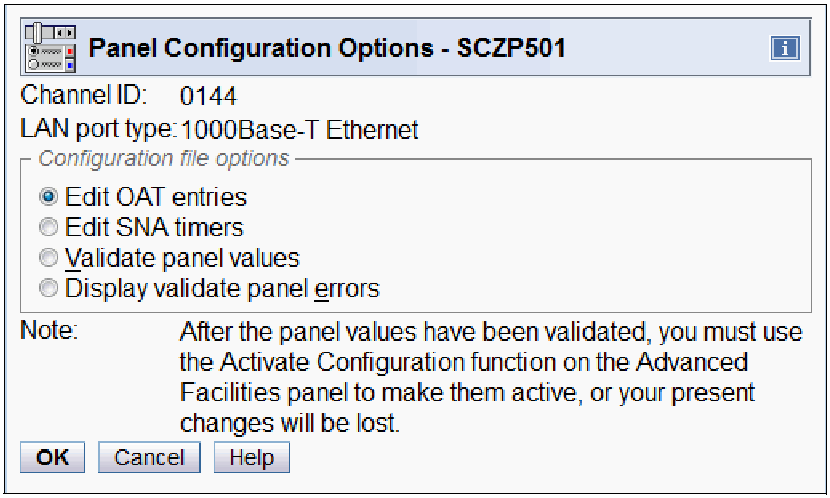

Figure 14-5 OSA/SF on HMC: Configuration file options

6. If you choose **Edit OAT entries**, the Edit OSA Address Table (OAT) Entries panel opens (Figure 14-6). You can edit OAT on this panel. CSS, Unit Address, and Device Number and so on are displayed based on your IOCDS configuration.

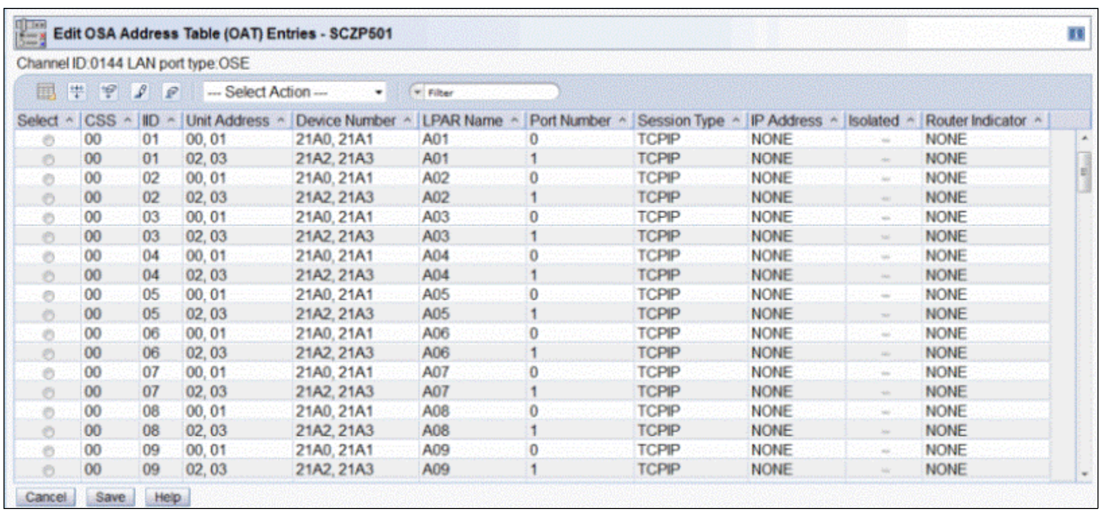

Figure 14-6 OSA/SF on HMC: Edit OSA Address Table (OAT) Entries panel

7. To edit, select the device in the left column (Figure 14-7). Then click **Select Action** and choose either **Edit as TCP/IP entry** or **Edit as SNA entry** for the selected device.

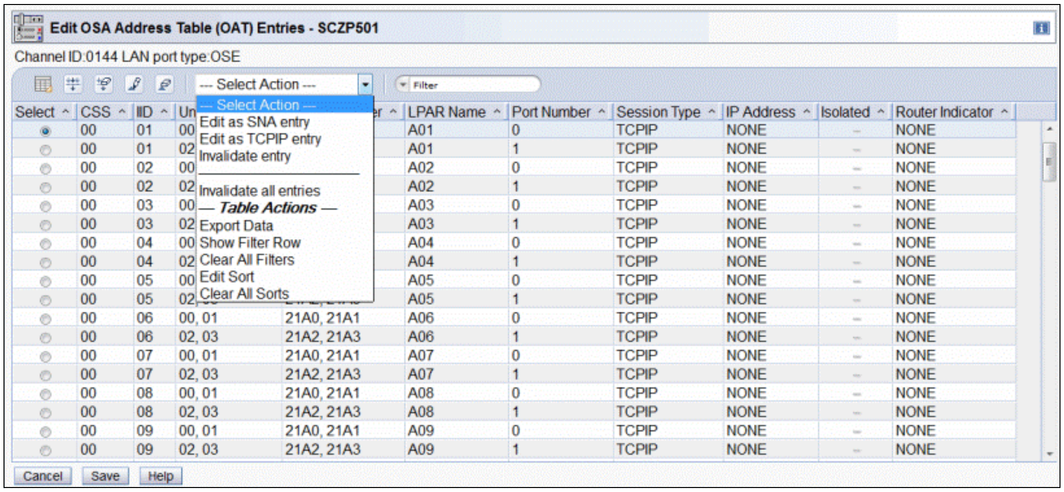

Figure 14-7 OSA/SF on HMC: Select Action for Edit OSA Address Table (OAT) Entries panel

– If you select **Edit as SNA Entry**, the OSA Address Table (OAT) Entry panel shown in Figure 14-8 opens. Select the appropriate SNA entry here and click **OK**.

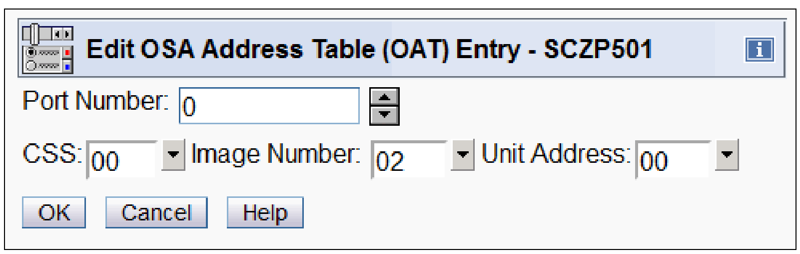

Figure 14-8 OSA/SF on HMC: Edit as SNA entry panel

– If you select **Edit as TCP/IP Entry**, the Edit OSA Address Table (OAT) Entry shown in Figure 14-9 on page 659 opens. Enter and select appropriate TCP/IP parameters here, and then click **OK**.

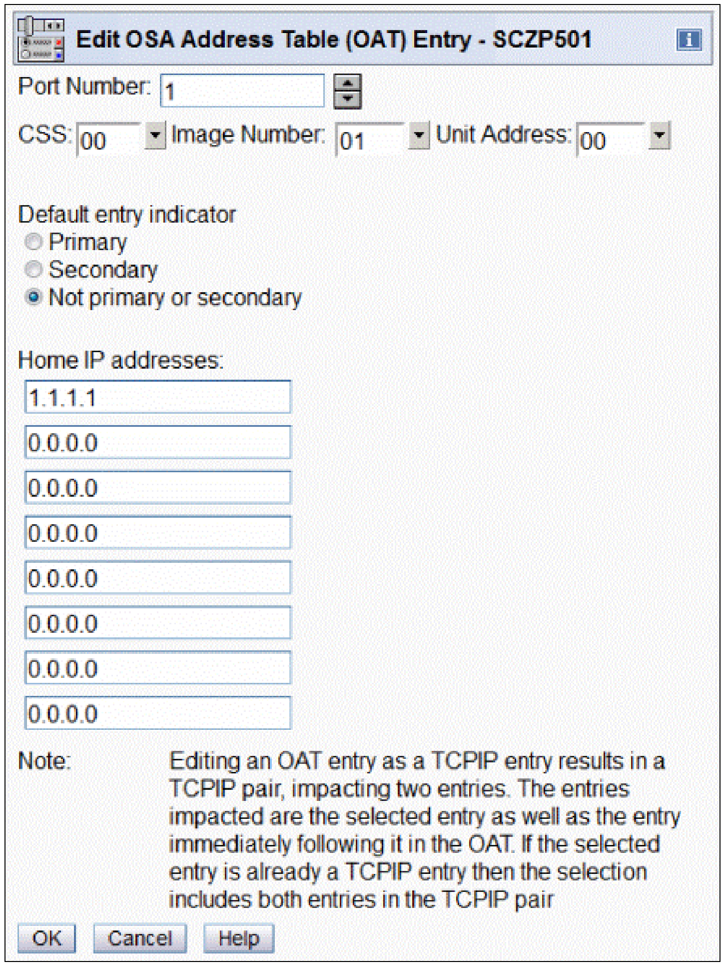

Figure 14-9 OSA/SF on HMC: Edit as TCP/IP entry panel

8. After editing TCP/IP entry or SNA entry, the Edit OSA Address Table (OAT) Entries panel opens again. Confirm that your entries are displayed on this panel. Figure 14-10 shows that the TCP/IP entry and SNA entry can be confirmed.

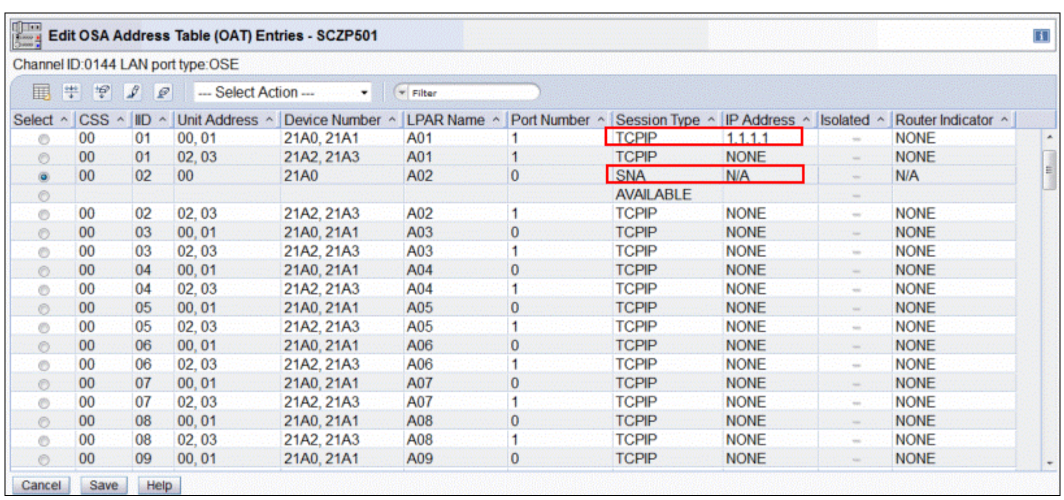

Figure 14-10 OSA/SF on HMC: Edit OSA Address Table (OAT) Entries panel after editing

9. If you select **Edit SNA timers** in the Panel Configuration Options panel (Figure 14-5 on page 657), you see the Edit the SNA Timers panel. If you change the SNA timer setting, edit the value on this panel (Figure 14-11).

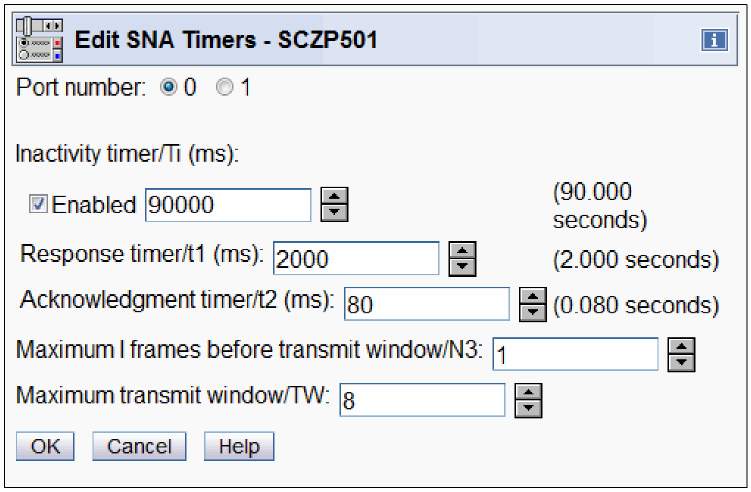

Figure 14-11 OSA/SF on HMC; Edit SNA Timers panel

10. When you complete editing the OAT entries or SNA timer, the Panel Configuration Options panel is displayed again. To activate your setting, you must validate your entries by selecting **Validate panel values**. Then, click **OK** (Figure 14-12).

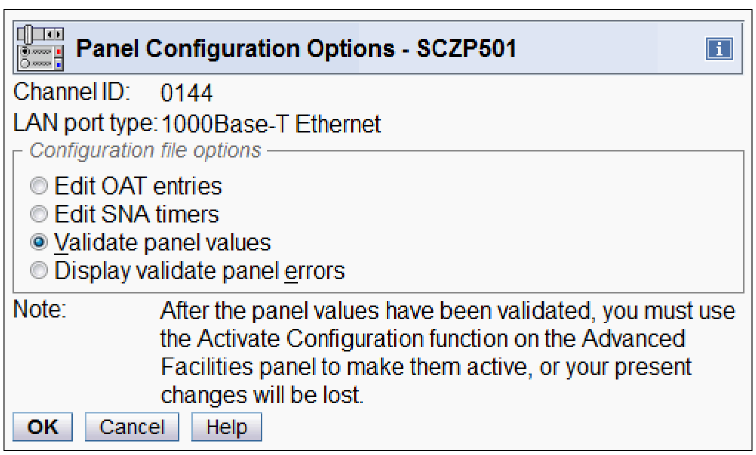

Figure 14-12 OSA/SF on HMC: Validate panel values

If your parameter has no error, a window similar to Figure 14-13 is displayed.

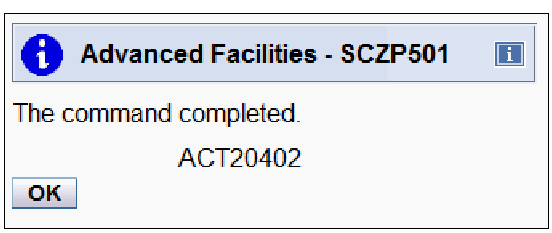

Figure 14-13 OSA/SF on HMC: Validate panel value with success

Otherwise, a window with ACT20425 message (Figure 14-14) opens, indicating that an error was found on the entry. To find the reason for the error, select **Display validate panel errors** (Figure 14-12 on page 660), and correct the error. Then validate the panel value again until the error is fixed.

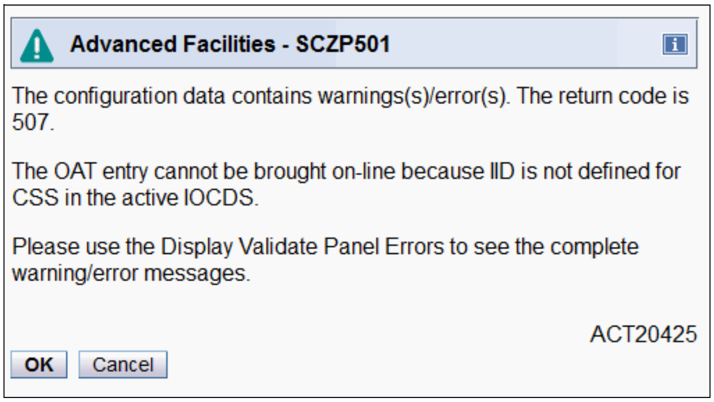

Figure 14-14 OSA/SF on HMC: error on Validate panel value

11. To activate your OSA configuration, select **Activate configuration** from the Advanced Facilities panel, and click **OK** (Figure 14-15).

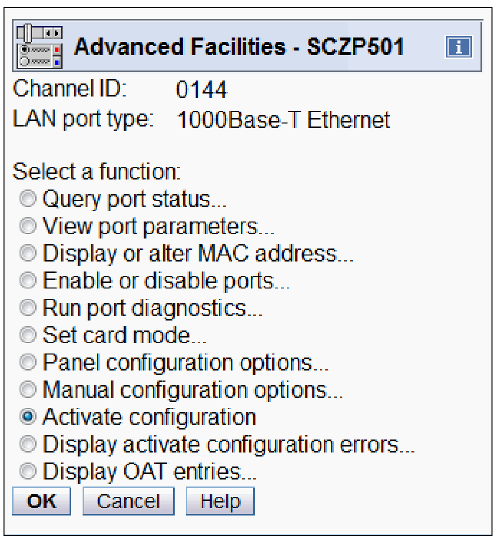

Figure 14-15 OSA/SF on HMC: Activate configuration

12. The Confirmation panel (Figure 14-16) opens. Click **Yes** to continue the operation.

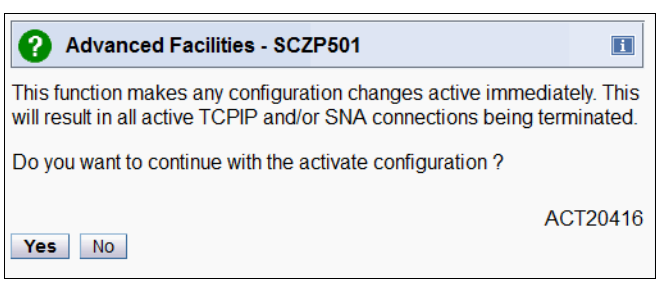

Figure 14-16 OSA/SF on HMC: Confirm activation

13.When the activation is successfully ended, the next panel opens (Figure 14-17). Click **OK** to complete the process.

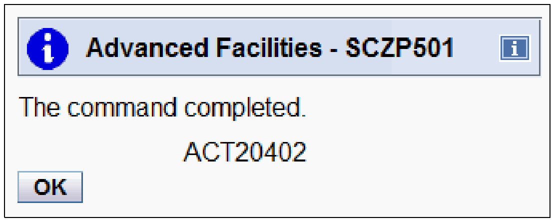

Figure 14-17 OSA/SF on HMC: message when activation completed

14.To reflect these customizations to OSA CHPID, *you must bring the CHPID OFFLINE from all the LPARs which share the OSA CHPID, and then back ONLINE*. Without this operation, you might have trouble with the OAT configuration.

15.If you select **Manual configuration options** (Figure 14-4 on page 656), the panel shown in Figure 14-18 opens. In this panel, you can import/export the source file of OAT through a USB device or FTP, create a new configuration file on the editor on HMC, and edit the source file. For details, see *Open Systems Adapter/Support Facility on the Hardware Management Console,* SC14-7580.

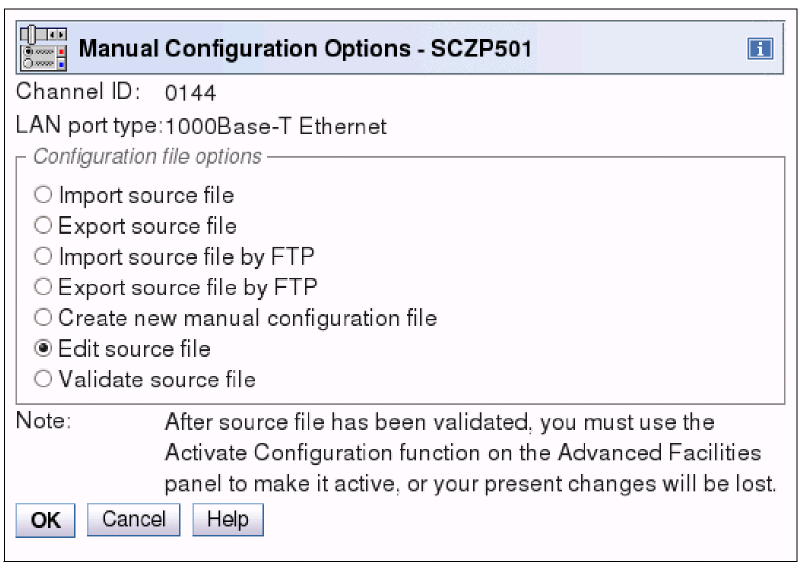

Figure 14-18 OSA/SF on HMC: Manual Configuration options

14.2 10GbE RoCE Express HCD definitions and z/OS display commands

In this section, we show how to define 10GbE RoCE Express SR (FC 0411) to z13 and we introduce the z/OS commands that are related to the PCIe features.

14.2.1 Understanding native PCIe feature definition rule

The 10GbE RoCE Express SR is a native PCIe I/O feature, so the HCD and IOCP definition rule differ from a non-PCIe card such as OSA Express-5S. Here are those rules:

- PCIe *Function Identifier (FID)* must be defined in HCD or HCM to create IOCP input.
  - FID is a hexadecimal value (three digits have the 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 has 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 function Identifier (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.

  Value examples include a RoCE feature that supports up to 31 partitions, and a zEDC Express feature that supports up to 15.
- Might have other parameters that are specific to the PCIe feature. For example, the 10GbE RoCE Express requires a Physical Network Identifier (PNETID=).
- 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. Methods are as follows:
  - Manually, by using the configurator (eCONFIG) PCHID report.
  - With assistance, by using the CHPID Mapping tool and with eConfig Configuration Report File (CFR) input.

> **Note:** Unlike CHPIDs, multiple functions can be mapped to the same PCHID. This is conceptually similar to mapping multiple InfiniBand coupling CHPIDs to the same adapter and port.

14.2.2 Consider Native PCIe feature Plugging and Resource Groups (RGs)

The native PCIe feature is provided by Resource Group (RG) code running on the system integrated firmware processor (IFP). For resilience, there are always two independent RGs on the system, sharing the IFP. For high availability purposes, always use two PCIe features located in different RGs (Figure 14-19).

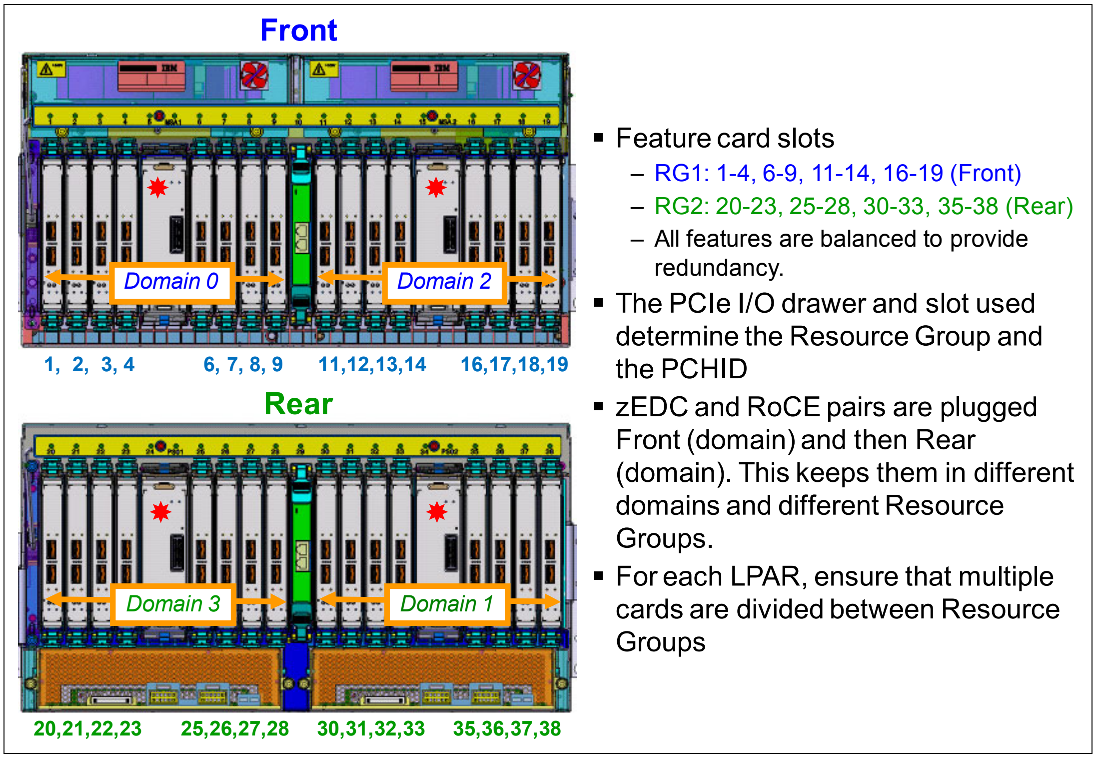

Figure 14-19 Relationship among PCIe I/O drawer slots, domains, and RGs in the z13

You can confirm which PCHID belongs to the RGs on the eConfig “AO” Data. eConfig places (plugs) hardware into slots to maximize serviceability and availability. Figure 14-20 is an example of eConfig output. You can see feature code RoCE Express (0411) located in RG1 and RG2 (Figure 14-20).

```
Source        Cage Slot F/C  PCHID/Ports or AID  Comment
A19/LG15/J01 Z22B  03  0420    108             RG1
A19/LG15/J01 Z22B  04  0411    10C/D1D2        RG1
A15/LG15/J01 Z22B  18  0411    138/D1D2        RG1
A15/LG15/J01 Z22B  19  0420    13C             RG1
A19/LG02/J01 Z22B  20  0420    140             RG2
A19/LG02/J01 Z22B  21  0411    144/D1D2        RG2
A15/LG02/J01 Z22B  37  0411    178/D1D2        RG2
A15/LG02/J01 Z22B  38  0420    17C             RG2
```

Figure 14-20 Sample eConfig PCHID report

Sample IOCP FUNCTION statements of RoCE Express are shown in Example 14-1.

Example 14-1 IOCP FUNCTION statements of RoCE Express

```
*10GbE RoCE Express Functions for LPAR LP14, Reconfigurable to LP03 or LP04
FUNCTION FID=9,VF=1,PART=((LP14),(LP03,LP04)),PNETID=(NET1,NET2), *
TYPE=ROCE,PCHID=10C
FUNCTION FID=A,VF=1,PART=((LP14),(LP03,LP04)),PNETID=(NET1,NET2), *
TYPE=ROCE,PCHID=144
*10GbE RoCE Express Functions for LPAR LP15, Reconfigurable to LP03 or LP04
FUNCTION FID=B,VF=2,PART=((LP15),(LP03,LP04)),PNETID=(NET1,NET2), *
TYPE=ROCE,PCHID=10C
FUNCTION FID=C,VF=2,PART=((LP15),(LP03,LP04)),PNETID=(NET1,NET2), *
TYPE=ROCE,PCHID=144
```

These are the key points of the sample IOCP FUNCTION.

- A function is identified by a function Identifier (FID): hexadecimal 000 – FFF.

Suggestion: Use *odd* FIDs for RG1 hardware; use *even* FIDs for RG2 hardware.

- TYPE is mandatory option to classify PCIe feature.
- Functions can be dedicated or reconfigurable, but not shared, only one LPAR in the Access List.
- To "share" a feature among LPARs, define a function for each LPAR. These functions must have different Function IDs (FIDs) and different Virtual Function ID (VF): decimal 1 - n is the maximum number LPARs the feature supports.
- Physical Network Identifier (PNETID) is positional to identify the network names: Port D1 (top), Port D2 (bottom). See Figure 14-20 on page 664.
- LP14 and LP15 have access to each type of hardware in both Resource Groups and, on RoCE, to both networks.

14.2.3 HCD definitions for RoCE Express

In this section, we describe how to configure RoCE Express in HCD. To define RoCE Express in HCD, these items must be defined:

- Function ID (FID)
- Type
- PCHID
- Virtual Function ID (VF)
- Physical Network Identifier (PNETID)

The Physical Network Identifier is a value that is defined to uniquely identify your physical layer 2 LAN fabric or physical broadcast domain. You can use this value to logically associate the z Systems features, adapters, and ports to be physically connected to your network. For more information about PNETID and how to configure RoCE Express in TCP/IP, see the information about Shared Memory Communications over Remote Direct Memory Access in *z/OS V2R1.0 Communications Server: IP Configuration Guide,* SC27-3650.

Complete the following steps:

1. Open HCD. On the Primary Task Selection panel, select **1. Define, modify, or view configuration data** (Figure 14-21).

```
z/OS V2.1 HCD
Command ===> _______________________________________________________

                           Hardware Configuration

Select one of the following.

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

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

I/O definition file . . . 'SYS1.IODF01.WORK'                    +
```

Figure 14-21   HCD: Primary Task

2. In the next panel (Define, Modify, or View Configuration Data), select **3. Processors** (Figure 14-22).

```
-------------- Define, Modify, or View Configuration Data ---------------

     Select type of objects to define, modify, or view data.

      3_ 1. Operating system configurations
              consoles
              system-defined generics
              EDTs
                esoterics
                user-modified generics
         2. Switches
              ports
              switch configurations
                port matrix
         3. Processors
              channel subsystems
                partitions
                channel paths
              PCIe functions
         4. Control units
         5. I/O devices
         6. Discovered new and changed control units and I/O devices
```

Figure 14-22   HCD: Define, Modify, or View Configuration Data

3. HCD displays the Processor List of defined processors. Type a forward slash (/) to select a Processor ID to define RoCE Express (Figure 14-23).

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

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


/ Proc. ID Type +   Model +  Mode+ Serial-# + Description
_ SCZP201  2097     E26      LPAR  01DE502097 Eclipse
_ SCZP401  2827     H43      LPAR  00B8D72827 Helix
/ SCZP501  2964     N63      LPAR  08DA872964 Sphinx
```

Figure 14-23 HCD: Processor List

4. In the Actions on selected processors panel, select **8. Work with PCIe functions (f)** (Figure 14-24).

.

```
Goto  Filter  Backup  Query  Help
------------- -------------- Actions on selected processors -----------------
              |                                                              |
Command ===>  |                                                              |
              | Select by number or action code and press Enter.             |
Select one or |                                                              |
              | 8_  1.  Add like . . . . . . . . . . . . . . . . (a)         |
              |     2.  Repeat (Copy) processor configurations (r)           |
/ Proc. ID Ty |     3.  Change . . . . . . . . . . . . . . . . . (c)         |
_ SCZP201  20 |     4.  *Prime serial number . . . . . . . . . (i)           |
_ SCZP401  28 |     5.  Delete . . . . . . . . . . . . . . . . . (d)         |
/ SCZP501  29 |     6.  View processor definition  . . . . . . (v)           |
************* |     7.  View related CTC connections . . . . . (k)           |
              |     8.  Work with PCIe functions . . . . . . . (f)           |
              |     9.  Work with partitions . . . . . . (SMP) (p)           |
              |     10. Work with attached channel paths (SMP) (s)           |
              |     11. Work with attached devices . . . (SMP) (u)           |
              |     12. Copy to channel subsystem  . . . (SMP) (y)           |
              |     13. Work with channel subsystems . . (XMP) (p,s)         |
              |                                                              |
              | * = requires TSA I/O Operations                              |
              |                                                              |
              ----------------------------------------------------------------
```

Figure 14-24 HCD: PCIe Actions on selected processors panel

5. In the PCIe Function List panel, you can define a new PCIe function by pressing F11 (Figure 14-25).

```
-------------------------------------------------------------------------------
                              PCIe Function List
Command ===> ______________________________________________ Scroll ===> PAGE

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

Processor ID . . . . : SCZP501      Sphinx



/ FID    PCHID  VF+  Type+           Description
******************************* Bottom of data ********************************
```

Figure 14-25 HCD: PCIe Function List

6. The Add PCIe Function panel opens (Figure 14-26). Define the following items and press Enter:
   – Function ID
   – Type
   – PCHID
   – Virtual Function ID

   In our example, we define the values that are shown in the figure.

```
----------------------------- Add PCIe Function ---------------------------


 Specify or revise the following values.

 Processor ID  . . . . : SCZP501      Sphinx

 Function ID . . . . . . 000
 Type  . . . . . . . . . ROCE          +

 PCHID . . . . . . . . . 140
 Virtual Function ID . . 1_  +

 Description . . . . . . RoCE Express ___________________
```

Figure 14-26 HCD: PCIe Add PCIe Function panel

7. The Add/Modify Physical Network IDs panel opens (Figure 14-27). On this panel, define the Physical Network ID (PNETID), which consists of a name with up to 16 alphanumeric characters. You can define two PNETIDs here. In this example, we define ITSOPNET as PNETID1. When you finish, press Enter.

```
--------------------- Add/Modify Physical Network IDs ---------------------


 If the PCHID is associated to one or more physical networks, specify
 each physical network ID corresponding to each applicable physical port.

 Physical network ID 1  . . ITSOPNET________
 Physical network ID 2  . . ________________
 Physical network ID 3  . . ________________
 Physical network ID 4  . . ________________
```

Figure 14-27   HCD: PCIe Add/Modify Physical Network IDs panel

8. The Define Access List panel opens (Figure 14-28). Specify the LPAR to be included in the Access List. In this example, we specify A01 in the access list.

Note: You can define only one LPAR in the Access List.

```
-------------------------- Define Access List ------------------------------
                                                               Row 1 of 19
 Command ===> _________________________________________ Scroll ===> PAGE

 Select one or more partitions for inclusion in the access list.

 Function ID  . . . . : 000

 / CSS ID Partition Name   Number Usage Description
 _ 0      A0A              A      OS
 _ 0      A0B              B      CF/OS CHPID holder
 _ 0      A0C              C      CF    CF7A
 _ 0      A0D              D      CF    CF7B
 _ 0      A0E              E      CF    CF7C
 _ 0      A0F              F      CF    CF7D
 / 0      A01              1      OS    SC76
 _ 0      A02              2      OS    SC74
 _ 0      A03              3      OS    SC75
 _ 0      A04              4      OS    ZZ0VM1
 _ 0      A05              5      OS    zAware1
 _ 0      A06              6      OS
 _ 0      A07              7      OS
 _ 0      A08              8      OS
```

Figure 14-28   HCD: PCIe Define Access List panel

9. The Define Candidate List panel opens (Figure 14-29). Define the LPARs to be included in the candidate list. in this example, we define A02 and A03 in the candidate list. Then press Enter to complete the operation.

```
-------------------------- Define Candidate List --------------------------
                                                               Row 1 of 18
 Command ===> _________________________________________ Scroll ===> PAGE

 Select one or more partitions for inclusion in the candidate list.

 Function ID  . . . . : 000

 / CSS ID Partition Name   Number Usage Description
 _ 0      A0A              A      OS
 _ 0      A0B              B      CF/OS CHPID holder
 _ 0      A0C              C      CF    CF7A
 _ 0      A0D              D      CF    CF7B
 _ 0      A0E              E      CF    CF7C
 _ 0      A0F              F      CF    CF7D
 / 0      A02              2      OS    SC74
 / 0      A03              3      OS    SC75
 _ 0      A04              4      OS    ZZ0VM1
 _ 0      A05              5      OS    zAware1
 _ 0      A06              6      OS
 _ 0      A07              7      OS
 _ 0      A08              8      OS
 _ 0      A09              9      OS
```

Figure 14-29 HCD: PCIe Define Candidate List

Now in the PCIe Function List, you see the FID that you defined (Figure 14-30).

```
------------------------------------------------------------------------------
                              PCIe Function List      Row 1 of 1 More:       >
Command ===> _______________________________________________ Scroll ===> PAGE

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

Processor ID . . . . : SCZP501      Sphinx



/ FID    PCHID  VF+  Type+           Description
_ 000    140    1    ROCE            RoCE Express
******************************* Bottom of data ********************************
```

Figure 14-30 HCD: PCIe Function List FID definition completed

The RoCE Express definitions that were just completed generates the IOCDS statements shown in Figure 14-31.

```
FUNCTION FID=000,VF=1,PCHID=140,PNETID=ITSOPNET1,              *
      PART=((A01),(A02,A03)),TYPE=ROCE
```

Figure 14-31 RoCE Express definition IOCDS deck

14.2.4 z/OS display commands for RoCE Express

In this section, we introduce the z/OS command related to the PCIe feature and show their response on our test system.

DISPLAY PCIE

You can use the **DISPLAY PCIE** command to display these items:

- All registered device drivers (with assigned printable names).
- All available or in-use PCIe functions and their associated device types.
- Information about a specific PCIe device with a list of the client address spaces using the device.

Figure 14-32 is an examples of the **D PCIE** command. You can confirm the FID and VFID you defined. FID is represented as PFID (PCIE function identifiers).

```
D PCIE
IQP022I 17.48.19 DISPLAY PCIE 600
PCIE     0010 ACTIVE
PFID  DEVICE TYPE NAME         STATUS  ASID  JOBNAME  PCHID VFN
0000  10GbE RoCE Express       CNFG                   0140  0001
```

Figure 14-32 example of D PCIE command

Figure 14-33 is an example of the **D PCIE,PFID=pfid** command. After you define the new PCIe function, enter this command and confirm its status is `ACTIVE`.

```
D PCIE,PFID=000
IQP024I 17.48.25 DISPLAY PCIE 602
PCIE     0010 ACTIVE
PFID  DEVICE TYPE NAME         STATUS  ASID  JOBNAME  PCHID VFN
0000  10GbE RoCE Express       CNFG                   0140  0001
 CLIENT ASIDS: NONE
```

Figure 14-33 example of D PCIE,PFID=pfid command

Figure 14-34 is example of the **D PCIE,DD** command. You can confirm the details of the device drives installed in the system.

```
D PCIE,DD
IQP023I 17.47.01 DISPLAY PCIE 596
PCIE     0010 ACTIVE
DEV TYPE   DEVICE TYPE NAME
1014044B   Hardware Accelerator
15B36750   10GbE RoCE
15B31003   10GbE RoCE
15B31004   10GbE RoCE Express
```

Figure 14-34 example of D PCIE,DD command

CONFIG command

You can use the **CONFIG** command to bring the PCIE function identifiers (PFID) online or offline.

Figure 14-35 is an example of **CONFIG PFID(x),ONLINE** command.

```
CF PFID(0),ONLINE
IEE504I PFID(0),ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Figure 14-35 example of CF PFID(x),ONLINE command

Figure 14-36 is an example of **CONFIG PFID(x),OFFLINE** command.

```
CF PFID(0),OFFLINE
IQP034I PCIE FUNCTION 00000000 OFFLINE. 652
 PCIE DEVICE TYPE NAME = (10GbE RoCE Express      ).
IEE505I PFID(0),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE
```

Figure 14-36 Example of CF PFID(x),OFFLINE command

14.3 HCD definitions and z/OS display commands for zBX

This section reviews the definitions in z/OS Hardware Configuration Definition (HCD) to support a zBX installation. For details about ZBX configuration, see Chapter 12, “Preparing IBM z13 for zBX Model 004” on page 595.

To support the internal communication networks that connect a z13 to a zBX, these CHPID types are used:

- CHPID type OSX for the Intraensemble Data Network (IEDN)
- CHPID type OSM for the Intranode Management Network (INMN)

For high availability purposes, you should define at least two OSM CHPIDs and two OSX CHPIDs.

14.3.1 Defining the OSM definitions

To define the OSM definitions, complete the following steps:

1. Make the necessary definitions to satisfy the INMN definitions that Ensemble management uses. In Figure 14-37, CHPID 0A is defined as type OSM. For this example, it is defined as spanned across multiple channel subsystems. Press Enter.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : SCZP501     Sphinx                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 0                                                  |
|                                                                           |
| Channel path ID  . . . . 0A    +           PCHID . . . 534                |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSM   +                                          |
| Operation mode . . . . . SPAN  +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________   +      |
| Description  . . . . . . Exp5s 1KBaseT                                    |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
+---------------------------------------------------------------------------+
```

Figure 14-37 HCD: OSM CHPID definition

2. The HCD prompt shown in Figure 14-38 opens. OSM channels must be defined to allow for more than 160 TCP/IP stacks. 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 14-38 HCD: OSM CHPID allowing 160 TCP/IP stacks

3. The HCD panel shown in Figure 14-39 on page 674 defines the *access list* for the OSM CHPID. LPARs with a forward slash (/) have access to CHPID 0A. In this example, there are no LPARs in the *candidate list*. Press Enter.

```
+--------------------------- Define Access List ----------------------------+
|                                                              Row 1 of 54 |
| Command ===> _________________________________________ Scroll ===> CSR   |
|                                                                           |
| Select one or more partitions for inclusion in the access list.          |
|                                                                           |
| Channel subsystem ID : 0                                                  |
| Channel path ID  . . : 0A     Channel path type  . : OSM                  |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| _ 0      A0A              A      OS    SC58                               |
| _ 0      A0B              B      OS    CHPID holder                       |
| _ 0      A0C              C      CF    WTSCPLX2 Sandbox CF1               |
| _ 0      A0D              D      CF    TESTPLEX CF7A                      |
| _ 0      A0E              E      CF    CF8A                               |
| _ 0      A0F              F      CF    CF8B                               |
| / 0      A01              1      OS    WTSCPLX8 SC80                      |
| / 0      A02              2      OS    VMLINUXA                           |
| / 0      A03              3      OS    WTSCPLX8 SC81                      |
| _ 0      A04              4      OS    SC76                               |
+---------------------------------------------------------------------------+
```

Figure 14-39 OSM CHPID access list

4. After the CHPID is defined, the Control Unit is defined as type OSM (Figure 14-40). Press Enter.

```
+------------------------- Add Control Unit --------------------------+
|                                                                      |
|                                                                      |
| Specify or revise the following values.                              |
|                                                                      |
| Control unit number  . . . . FA40  +                                 |
| Control unit type  . . . . . OSM           +                         |
|                                                                      |
| Serial number  . . . . . . . __________                              |
| Description  . . . . . . . . OSM for INMN Connection                 |
|                                                                      |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +              |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +              |
|                                                                      |
| If connected to a switch:                                            |
|                                                                      |
| Define more than eight ports . . 2   1.  Yes                         |
|                                      2.  No                          |
| Propose CHPID/link addresses and                                     |
| unit addresses . . . . . . . . . 2   1.  Yes                         |
|                                      2.  No                          |
|  F1=Help    F2=Split   F3=Exit    F4=Prompt  F5=Reset   F9=Swap      |
| F12=Cancel                                                           |
+----------------------------------------------------------------------+
```

Figure 14-40 OSM Control Unit definition

5. Figure 14-41 shows defining the control unit address range. By default, the range is 255 addresses that start with 00. Press Enter.

```
+---------------------------- Add Control Unit -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Control unit number  . : FA40        Type . . . . . . : OSM               |
| Processor ID . . . . . : SCZP501     Sphinx                               |
| Channel Subsystem ID . : 0                                                |
|                                                                           |
| Channel path IDs . . . . 0A    __    __    __    __    __    __    __   + |
| 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 14-41 HCD: OSM Control Unit address range

6. Define the OSM devices (type OSA-M). In this example, we define 16 devices FA40-FA4F (Figure 14-42). Press Enter.

```
+------------------------------- Add Device ---------------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Device number  . . . . . . . . FA40  + (0000 - FFFF)                      |
| Number of devices  . . . . . . 16__                                       |
| Device type  . . . . . . . . . OSA-M________  +                           |
|                                                                           |
| Serial number  . . . . . . . . __________                                 |
| Description  . . . . . . . . . OSM for INMN connection________            |
|                                                                           |
| Volume serial number . . . . . ______  (for DASD)                         |
|                                                                           |
| PPRC usage . . . . . . . . . . _  + (for DASD)                            |
|                                                                           |
| Connected to CUs . . FA40  ____  ____  ____  ____  ____  ____  ____  +    |
|                                                                           |
+---------------------------------------------------------------------------+
```

Figure 14-42 HCD: OSM device add

7. After you select the processor ID, press Enter. The window shown in Figure 14-43 opens.

```
+------------------------- Define Device / Processor -------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number  . . . : FA40          Number of devices . . . . : 16         |
| Device type  . . . . : OSA-M                                                |
| Processor ID . . . . : SCZP501      Sphinx                                  |
| Channel Subsystem ID : 0                                                    |
|                                                                             |
| Subchannel set ID  . . . . . . . _   +                                      |
| Unit address . . . . . . . . . . 40  + (Only necessary when different from |
|                                      the last 2 digits of device number)    |
| Time-Out . . . . . . . . . . . . No   (Yes or No)                           |
| STADET . . . . . . . . . . . . . No   (Yes or No)                           |
|                                                                             |
| Preferred CHPID  . . . . . . . . __  +                                      |
| Explicit device candidate list . No   (Yes or No)                           |
|                                                                             |
+-----------------------------------------------------------------------------+
```

Figure 14-43 HCD: OSM device add processor definition

8. After the subchannel device is added, the devices must be added to the appropriate Operating System Configurations, as shown in Figure 14-44. Press Enter.

```
+----------- Define Device to Operating System Configuration -----------+
|                                                            Row 1 of 5 |
| Command ===> _____________________________________ Scroll ===> CSR    |
|                                                                       |
| Select OSs to connect or disconnect devices, then press Enter.        |
|                                                                       |
| Device number  . : FA40          Number of devices  : 16              |
| Device type  . . : OSA-M                                              |
|                                                                       |
| / Config. ID   Type     SS Description                      Defined   |
| _ ALLDEV       MVS         All devices                                |
| _ LABSERV1     MVS         Lab Services                               |
| / L06RMVS1     MVS         Sysplex systems                            |
| _ MVSW1        MVS         Production systems                         |
| _ TRAINER      MVS         Trainer - Local Site Online                |
| ************************** Bottom of data *************************** |
+-----------------------------------------------------------------------+
```

Figure 14-44 HCD: OSM device adding OS Config

The panel shown in Figure 14-45 determines the device parameters for the specific operating system configuration.

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                    Row 1 of 3 |
| Command ===> ___________________________________________ Scroll ===> CSR     |
|                                                                               |
| Specify or revise the values below.                                           |
|                                                                               |
| Configuration ID . : L06RMVS1     Sysplex systems                             |
| Device number  . . : FA40         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 14-45 HCD: Device OS parameters

The OSM definitions that were just completed generate the IOCDS statements shown in Figure 14-46.

For OSM channel paths, device priority queuing must be disabled. Therefore, for these CHPIDs, HCD converts a CHPARM=00 (defaults to priority queuing enabled) to `CHPARM=02`. The statements generated depict the CHPID `Type=OSM`, the Control Unit Type (`UNIT=OSM`), and the IODEVICE type (`UNIT=OSA`) with `MODEL-M` parameter.

```
CHPID PATH=(CSS(0,1,2),0A),SHARED,                             *
      PARTITION=((CSS(0),(A01,A02,A03,A07),(=)),(CSS(1),(A11,A*
      12,A13,A16,A17,A18,A19),(=)),(CSS(2),(A25),(=))),        *
      CHPARM=02,PCHID=534,TYPE=OSM
CNTLUNIT CUNUMBR=FA40,                                         *
      PATH=((CSS(0),0A),(CSS(1),0A),(CSS(2),0A)),UNIT=OSM
IODEVICE ADDRESS=(FA40,016),MODEL=M,UNITADD=00,CUNUMBR=(FA40),*
      UNIT=OSA
```

Figure 14-46 HCD: OSM IOCDS statements

14.3.2 Defining the OSX definitions

The definitions described in this section are necessary to satisfy the IEDN requirements.

In Figure 14-47, the CHPID 18 is defined as type OSX. In this example, the CHPID is defined as SPANNED.

```
+---------------------------- Add Channel Path -----------------------------+
|                                                                           |
|                                                                           |
| Specify or revise the following values.                                   |
|                                                                           |
| Processor ID . . . . : SCZP501     Sphinx                                 |
| Configuration mode . : LPAR                                               |
| Channel Subsystem ID : 0                                                  |
|                                                                           |
| Channel path ID  . . . . 18    +            PCHID . . . 538               |
| Number of CHPIDs . . . . 1                                                |
| Channel path type  . . . OSX   +                                          |
| Operation mode . . . . . SPAN  +                                          |
| Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +       |
| Description  . . . . . . Exp5S 10GbE SR                                   |
|                                                                           |
| Specify the following values only if connected to a switch:               |
| Dynamic entry switch ID  __  + (00 - FF)                                  |
| Entry switch ID  . . . . __  +                                            |
| Entry port . . . . . . . __  +                                            |
```

Figure 14-47 HCD: OSX CHPID definition

If you require more than 160 TCP/IP stacks, reply `YES` to the prompt. In this case, more than 160 stacks are not required, so we accept the default answer of `No`, as shown in Figure 14-48.

```
+----------- Allow for more than 160 TCP/IP stacks -----------+
|                                                             |
|                                                             |
| Specify Yes to allow more than 160 TCP/IP stacks,           |
| otherwise specify No. Specifying Yes will cause priority    |
| queuing to be disabled.                                     |
|                                                             |
| Will greater than 160 TCP/IP stacks                         |
| be required for this channel?  . . . No                     |
+-------------------------------------------------------------+
```

Figure 14-48 HCD: OSX CHPID not allowing more than 160 TCP/IP stacks

The CHPID access list is defined (Figure 14-49) for the OSX CHPID. The LPARs selected with a forward slash (/) have access to the CHPID. In this example, there are no LPARs in the candidate list, so this panel is skipped.

```
+--------------------------- Define Access List ----------------------------+
|                                                              Row 1 of 54 |
| Command ===> _________________________________________ Scroll ===> CSR    |
|                                                                           |
| Select one or more partitions for inclusion in the access list.           |
|                                                                           |
| Channel subsystem ID : 0                                                  |
| Channel path ID  . . : 18     Channel path type  . : OSX                  |
| Operation mode . . . : SPAN   Number of CHPIDs . . : 1                    |
|                                                                           |
| / CSS ID Partition Name   Number Usage Description                        |
| _ 0      A0A              A      OS    SC58                               |
| _ 0      A0B              B      OS    CHPID holder                       |
| _ 0      A0C              C      CF    WTSCPLX2 Sandbox CF1               |
| _ 0      A0D              D      CF    TESTPLEX CF7A                      |
| _ 0      A0E              E      CF    CF8A                               |
| _ 0      A0F              F      CF    CF8B                               |
| / 0      A01              1      OS    WTSCPLX8 SC80                      |
| / 0      A02              2      OS    VMLINUXA                           |
| / 0      A03              3      OS    WTSCPLX8 SC81                      |
| _ 0      A04              4      OS    SC76                               |
```

Figure 14-49 HCD: OSX CHPID access list

The OSX control unit is defined as type OSX (Figure 14-50).

```
+-------------------------- Add Control Unit ---------------------------+
|                                                                        |
|                                                                        |
| Specify or revise the following values.                                |
|                                                                        |
| Control unit number  . . . . 2300  +                                   |
| Control unit type  . . . . . OSX__________  +                          |
|                                                                        |
| Serial number  . . . . . . . __________                                |
| Description  . . . . . . . . IEDN connection_________________          |
|                                                                        |
| Connected to switches  . . . __ __ __ __ __ __ __ __  +                |
| Ports  . . . . . . . . . . . __ __ __ __ __ __ __ __  +                |
|                                                                        |
| If connected to a switch:                                              |
|                                                                        |
| Define more than eight ports . . 2   1.  Yes                           |
|                                      2.  No                            |
| Propose CHPID/link addresses and                                       |
| unit addresses . . . . . . . . . 2   1.  Yes                           |
|                                      2.  No                            |
```

Figure 14-50 HCD: OSX add Control Unit

The control unit address range is defined (Figure 14-51), which starts at 00 for 255 devices.

```
+----------------------------- Add Control Unit ------------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Control unit number  . : 2300        Type . . . . . . : OSX                 |
| Processor ID . . . . . : SCZP501     Sphinx                                 |
| Channel Subsystem ID . : 0                                                  |
|                                                                             |
| Channel path IDs . . . . 18    __    __    __    __    __    __    __    +  |
| 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 14-51 HCD: OSX add Control Unit for Processor

The OSX devices are defined (Figure 14-52) as type OSA-X with a device address range of 2300 - 230F.

```
+-------------------------------- Add Device ---------------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number  . . . . . . . . 2300  + (0000 - FFFF)                        |
| Number of devices  . . . . . . 16                                           |
| Device type  . . . . . . . . . OSA-X         +                              |
|                                                                             |
| Serial number  . . . . . . . . __________                                   |
| Description  . . . . . . . . . IEDN connection                              |
|                                                                             |
| Volume serial number . . . . . ______  (for DASD)                           |
|                                                                             |
| PPRC usage . . . . . . . . . . _  + (for DASD)                              |
|                                                                             |
| Connected to CUs . . 2300  ____  ____  ____  ____  ____  ____  ____  +      |
|                                                                             |
```

Figure 14-52 HCD: OSX Device definition

After you select the processor ID, the next panel opens (Figure 14-53). In the example, a unit address of 00 is selected.

```
+------------------------- Define Device / Processor -------------------------+
|                                                                             |
|                                                                             |
| Specify or revise the following values.                                     |
|                                                                             |
| Device number  . . . : 2300          Number of devices . . . . : 16         |
| Device type  . . . . : OSA-X                                                |
| Processor ID . . . . : SCZP501       Sphinx                                 |
| Channel Subsystem ID : 0                                                    |
|                                                                             |
| Subchannel set ID  . . . . . . . _   +                                      |
| Unit address . . . . . . . . . . 00  + (Only necessary when different from  |
|                                      the last 2 digits of device number)    |
| Time-Out . . . . . . . . . . . . No   (Yes or No)                           |
| STADET . . . . . . . . . . . . . No   (Yes or No)                           |
|                                                                             |
| Preferred CHPID  . . . . . . . . __  +                                      |
| Explicit device candidate list . No   (Yes or No)                           |
|                                                                             |
```

Figure 14-53 HCD: OSX device definition for processor

The devices are added to the L06RMVS1 operating system configuration (Figure 14-54).

```
+----------- Define Device to Operating System Configuration -----------+
|                                                           Row 1 of 5 |
| Command ===> _____________________________________ Scroll ===> CSR   |
|                                                                      |
| Select OSs to connect or disconnect devices, then press Enter.       |
|                                                                      |
| Device number  . : 2300          Number of devices  : 16             |
| Device type  . . : OSA-X                                             |
|                                                                      |
| / Config. ID   Type     SS Description                      Defined  |
| _ ALLDEV       MVS         All devices                                |
| _ LABSERV1     MVS         Lab Services                               |
| / L06RMVS1     MVS         Sysplex systems                            |
| _ MVSW1        MVS         Production systems                         |
| _ TRAINER      MVS         Trainer - Local Site Online               |
| ************************** Bottom of data *************************** |
```

Figure 14-54 HCD: OSX device being added to the operating system

Finally, the operating system configuration parameters are defined (Figure 14-55).

```
+-------------------- Define Device Parameters / Features --------------------+
|                                                                  Row 1 of 3 |
| Command ===> ___________________________________________ Scroll ===> CSR    |
|                                                                             |
| Specify or revise the values below.                                         |
|                                                                             |
| Configuration ID . : L06RMVS1     Sysplex systems                           |
| Device number  . . : 2300         Number of devices  : 16                   |
| Device type  . . . : OSA-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              |
| ***************************** Bottom of data ****************************** |
```

Figure 14-55 HCD: OSX device OS parameters

The IOCDS statements shown in Figure 14-56 are generated. Note the CHPID type of `OSX`, the control unit type of `OSX`, and the IODEVICE is `MODEL=X` and `UNIT=OSA`.

```
CHPID PATH=(CSS(0,1,2),18),SHARED,                                     *
      PARTITION=((CSS(0),(A01,A02,A03,A07),(=)),(CSS(1),(A11,A*
      12,A13,A16,A17,A18,A19),(=)),(CSS(2),(A25),(=))),        *
      PCHID=538,TYPE=OSX
CNTLUNIT CUNUMBR=2300,                                                 *
      PATH=((CSS(0),18),(CSS(1),18),(CSS(2),18)),UNIT=OSX
IODEVICE ADDRESS=(2300,016),MODEL=X,CUNUMBR=(2300),UNIT=OSA
```

Figure 14-56 OSM IOCDS statements

14.4 z/OS display commands

This section contains several common device status z/OS operator commands to display the status of the OSM and OSX type devices.

14.4.1 z/OS display commands for CHPID type OSM

The **D M=CHP(0A)** command shown in Figure 14-57 displays the CHPID path status for the OSM type 31 CHPID.

```
D M=CHP(0A)

IEE174I 12.52.44 DISPLAY M 848
CHPID 0A:    TYPE=31, DESC=OSA ZBX MANAGEMENT, ONLINE
DEVICE STATUS FOR CHANNEL PATH 0A
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
FA40 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 0534
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE      @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE     $ PATH NOT OPERATIONAL
```

Figure 14-57 OSM D M=CHP)

The **D M=DEV(FA40)** command and response are in Figure 14-58. The response displays the OSM device path status. Notice that the node descriptor information that is returned includes the pseudo control unit 1730.007 and 1732.001 used for the OSA-Express5S port. Also included is the 2964 machine type, serial number, and a TAG of 0A00 indicating the CHPID.

```
D M=DEV(FA40)

IEE174I 12.53.29 DISPLAY M 850
DEVICE 0FA40     STATUS=ONLINE
CHP                     0A
ENTRY LINK ADDRESS      ..
DEST LINK ADDRESS       0D
PATH ONLINE             Y
CHP PHYSICALLY ONLINE   Y
PATH OPERATIONAL        Y
MANAGED                 N
CU NUMBER               FA40
MAXIMUM MANAGED CHPID(S) ALLOWED:    0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 001730.007.IBM.02.29640008DA87.0A00
SCP TOKEN NED     = 001730.007.IBM.02.29640008DA87.0A00
SCP DEVICE NED    = 001732.001.IBM.02.29640008DA87.0A00
```

Figure 14-58 OSM D M=DEV(FA40)

The **D U,,,FA40,4** command and response are shown in Figure 14-59. The response shows the status for the OSM devices

```
D U,,,FA40,4

IEE457I 12.54.09 UNIT STATUS 854
UNIT TYPE STATUS          VOLSER     VOLSTATE
FA40 OSAM A-BSY
FA41 OSAM A-RAL
FA42 OSAM A-BSY
FA43 OSAM A-BSY
```

Figure 14-59 OSM D U,,,device

The **D U,,ALLOC,FA40,4** command and response are shown in Figure 14-60. The OSM device job allocation is shown as NET, which is the IBM VTAM® task.

```
D U,,ALLOC,FA40,4

IEE106I 12.54.57 UNITS ALLOCATED 858
UNIT      JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID
FA40      NET      001F
FA41      NET      001F
FA42      NET      001F
FA43      NET      001F
```

Figure 14-60 OSM D U,,alloc,device

14.4.2 z/OS display commands for CHPID type OSX

The **D M=CHP(18)** command and response are shown in Figure 14-61. The response shows the CHPID path status for the OSX type 30 CHPID.

```
D M=CHP(18)

IEE174I 12.55.24 DISPLAY M 865
CHPID 18:    TYPE=30, DESC=OSA ZBX DATA, ONLINE
DEVICE STATUS FOR CHANNEL PATH 18
     0  1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
2300 +  +  +  +  +  +  +  +  +  +  +  +  +  +  +  +
SWITCH DEVICE NUMBER = NONE
PHYSICAL CHANNEL ID = 0538
************************ SYMBOL EXPLANATIONS ************************
+ ONLINE      @ PATH NOT VALIDATED   - OFFLINE    . DOES NOT EXIST
* PHYSICALLY ONLINE     $ PATH NOT OPERATIONAL
```

Figure 14-61 OSX D M=CHP

The **D M=DEV(2300)** command and response are shown in Figure 14-62, indicating the OSX device path status. Notice that the node descriptor information that is returned includes the pseudo control unit 1730.007 and 1732.001 used for the OSA-Express5S port. Also included is the 2964 machine type, serial number, and a TAG of 1000 indicating the CHPID.

```
D M=DEV(2300)

IEE174I 12.55.40 DISPLAY M 867
DEVICE 02300     STATUS=ONLINE
CHP                    18
ENTRY LINK ADDRESS     ..
DEST LINK ADDRESS      0D
PATH ONLINE            Y
CHP PHYSICALLY ONLINE  Y
PATH OPERATIONAL       Y
MANAGED                N
CU NUMBER              2300
MAXIMUM MANAGED CHPID(S) ALLOWED:    0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 001730.007.IBM.02.29640008DA87.1000
SCP TOKEN NED     = 001730.007.IBM.02.29640008DA87.1000
SCP DEVICE NED    = 001732.001.IBM.02.29640008DA87.1000
```

Figure 14-62 OSX D M=DEV

The **D U,,,2300,4** command and response are shown in Figure 14-63. The response displays the status for the OSX devices.

```
D U,,,2300,4

IEE457I 12.55.53 UNIT STATUS 869
UNIT TYPE STATUS          VOLSER     VOLSTATE
2300 OSAX A-BSY
2301 OSAX A
2302 OSAX A-BSY
2303 OSAX A-BSY
```

Figure 14-63 OSX D U,,,device

The **D U,,ALLOC,2300,4** command and response are shown in Figure 14-64. The OSX device job allocation is NET, which is the VTAM task.

```
D U,,ALLOC,2300,4

IEE106I 12.56.02 UNITS ALLOCATED 871
UNIT      JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID
2300      NET      001F
2301      NET      001F
2302      NET      001F
2303      NET      001F
```

Figure 14-64 OSX D U,,alloc,device

14.5 Configure OSA-Express Integrated Console Controller (OSA-ICC)

When you define OSA-Express as OSC, it will be the controller for TN3270, operator console, and printer session.

If you upgrade from z196 or zEC12, your OSA-ICC configuration is automatically migrated to z13. You must configure OSA-ICC when you newly install z13 or replace from another z Systems Processor.

This section shows how to configure, migrate, and activate OSA-ICC configuration.

14.5.1 Defining the OSC definitions

Before you use OSA-ICC, define OSC CHPID, CNTLUNIT, and IODEVICE on IOCDS. Here is an example of how to define OSC in HCD:

1. Define the channel path to the processor. To use OSA-ICC, you must define the channel path type as `OSC` (Figure 14-65). If you use this OSC channel among more than one LCSS, set the operation mode as `SPAN`. Press Enter to continue.

```
---------------------------- Add Channel Path ----------------------------

 Specify or revise the following values.

 Processor ID . . . . : SCZP501      Sphinx
 Configuration mode . : LPAR
 Channel Subsystem ID : 0

 Channel path ID  . . . . 01 +           Channel ID  210  +
 Number of CHPIDs . . . . 1
 Channel path type  . . . OSC   +
 Operation mode . . . . . SPAN  +
 Managed  . . . . . . . . No   (Yes or No)   I/O Cluster ________  +
 Description  . . . . . . OSA Express5S OSA-ICC_________

 Specify the following values only if connected to a switch:
 Dynamic entry switch ID  __  + (00 - FF)
 Entry switch ID  . . . . __  +
 Entry port . . . . . . . __  +
```

Figure 14-65 CD: OSC Channel Path definition

2. The Define Access List panel opens(Figure 14-66). The LPARs that are selected with a forward slash (/) have access to the CHPID. Press Enter to continue.

```
------------------------- Define Access List ----------------------------
                                                              Row 1 of 19
Command ===> _________________________________________ Scroll ===> PAGE

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 0
Channel path ID  . . : 01     Channel path type  . : OSC
Operation mode . . . : SPAN   Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
/ 0      A0A              A      OS
_ 0      A0B              B      CF/OS CHPID holder
_ 0      A0C              C      CF    CF7A
_ 0      A0D              D      CF    CF7B
_ 0      A0E              E      CF    CF7C
_ 0      A0F              F      CF    CF7D
/ 0      A01              1      OS    SC76
/ 0      A02              2      OS    SC74
/ 0      A03              3      OS    SC75
/ 0      A04              4      OS    ZZ0VM1
_ 0      A05              5      OS    zAware1
```

Figure 14-66 HCD: Define Access List for OSC

3. The Define Candidate List panel (Figure 14-67) opens. If you want to include an LPAR in the candidate list, select it with a forward slash (/). In this example, we do not define a candidate list.

```
------------------------ Define Candidate List -------------------------
                                                               Row 1 of 8
Command ===> _________________________________________ Scroll ===> PAGE

Select one or more partitions for inclusion in the candidate list.

Channel subsystem ID : 0
Channel path ID  . . : 01     Channel path type  . : OSC
Operation mode . . . : SPAN   Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
_ 0      A0B              B      CF/OS CHPID holder
_ 0      A0C              C      CF    CF7A
_ 0      A0D              D      CF    CF7B
_ 0      A0E              E      CF    CF7C
_ 0      A0F              F      CF    CF7D
_ 0      A05              5      OS    zAware1
_ 1      A1B              B      CF/OS CHPID holder
_ 1      A1C              C      CF    WTSCPLX8 CF8C
```

Figure 14-67 HCD: Define Candidate List for OSC

4. Define the control unit type (CNTLUNIT) as `OSC` (Figure 14-68 and Figure 14-69).

```
------------------------ Add Control Unit -------------------------


Specify or revise the following values.

Control unit number  . . . . F401  +
Control unit type  . . . . . OSC__________   +

Serial number  . . . . . . . __________
Description  . . . . . . . . ________________________________

Connected to switches  . . . __  __  __  __  __  __  __  __   +
Ports  . . . . . . . . . . . __  __  __  __  __  __  __  __   +

If connected to a switch:

Define more than eight ports . . 2   1.  Yes
                                     2.  No
Propose CHPID/link addresses and
unit addresses . . . . . . . . . 2   1.  Yes
                                     2.  No
```

Figure 14-68 HCD: OSC Add Control Unit (part 1 of 2)

```
---------------------------- Add Control Unit ------------------------------


 Specify or revise the following values.

 Control unit number  . : F401        Type . . . . . . : OSC
 Processor ID . . . . . : SCZP501     Sphinx
 Channel Subsystem ID . : 0

 Channel path IDs . . . . 01    __    __    __    __    __    __    __     +
 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 14-69 HCD: OSC Add Control Unit (part 2 of 2)

5. Define IODEVICE as follows:

   a. Define the device type as 3270-X (Figure 14-70).

```
------------------------------ Add Device ---------------------------------


Specify or revise the following values.

Device number  . . . . . . . . F400  + (0000 - FFFF)
Number of devices  . . . . . . 1
Device type  . . . . . . . . . 3270-X         +

Serial number  . . . . . . . . __________
Description  . . . . . . . . . ________________________________

Volume serial number . . . . . ______  (for DASD)

PPRC usage . . . . . . . . . . _  + (for DASD)

Connected to CUs . . F400  ____  ____  ____  ____  ____  ____  ____  +
```

Figure 14-70   HCD: Add Device for OSC (part 1 of 3)

   b. Associate this device to OSCONFIG and define the appropriate parameter (Figure 14-71).

```
---------- Define Device to Operating System Configuration -----------
                                                           Row 1 of 1
 Command ===> _____________________________________ Scroll ===> PAGE

 Select OSs to connect or disconnect devices, then press Enter.

 Device number  . : F400          Number of devices  : 1
 Device type  . . : 3270-X

 / Config. ID   Type     SS Description                     Defined
 / L06RMVS1     MVS         Sysplex systems
 ************************** Bottom of data ***************************
```

Figure 14-71   HCD: Add Device for OSC (part 1 of 3)

   c. If you use this device as a NIP console, then it must be defined in the NIP Console List (see Figure 14-72 on page 690).

```
-------------------------- NIP Console List -------------------------
   Goto  Backup  Query  Help
---------------------------------------------------------------------

 Command ===> _______________________________________ Scroll ===> PAGE

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

 Configuration ID . : L06RMVS1     Sysplex systems
                  -------------- Add NIP Console ----------------
   Order   Device |                                              |
 / Number  Number |                                              |
 **************** | Specify the following values.                | *****
                  |                                              |
                  | Device number of console . . . . . . F400    |
                  |                                              |
                  | Order number . . . . . . . . . . . . 1       |
                  |                                              |
                  |                                              |
                  |                                              |
                  ------------------------------------------------
```

Figure 14-72 HCD: Add Device for OSC (part 1 of 3)

The completed OSC definitions generate the IOCDS statements shown in Figure 14-73.

```
CHPID PATH=(CSS(0),01),SHARED,                                      *
      PARTITION=((A0A,A01,A02,A03,A04,A06,A07,A08,A09),(=)),  *
      PCHID=210,TYPE=OSC
CNTLUNIT CUNUMBR=F401,PATH=((CSS(0),01)),UNIT=OSC
IODEVICE ADDRESS=(F400,001),MODEL=X,CUNUMBR=(F401),UNIT=3270
```

Figure 14-73 OSA-ICC IOCDS deck sample

14.5.2 Saving and restoring OSA-ICC configuration

When you replace z13, you can export an OSA-ICC configuration file from another z Systems platform and export it to z13. In this section, we describe how to export and import the OSA-ICC configuration file using HMC.

Exporting OSA-ICC configuration file using OSA Advanced Facilities

In this example, we export the OSA-ICC configuration file from zEC12 (the CPC name is SCZP401) and import the file to z13 (the CPC name is SCZP501).

To export the OSA-ICC configuration file, follow these 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 panel shown in Figure 14-74 is displayed.

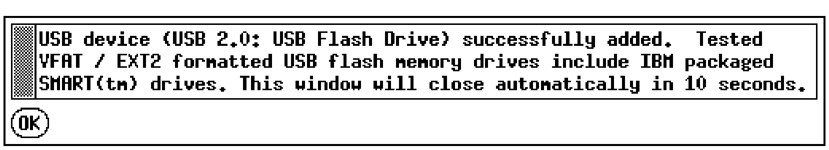

Figure 14-74 OSA Advanced Facilities: removable media inserted

2. Log on to the HMC, select the CPC you want to operate, and open the OSA Advanced facility.
3. Select OSC CHPID to export the OSA-ICC configuration file. Then, select **Card specific advanced facilities.** This action is same as is shown in step 1 on page 655.
4. Select **Manual configuration options** and click **OK.**
5. The Manual Configuration Options panel opens (Figure 14-75). Select **Export source file** and click **OK**.

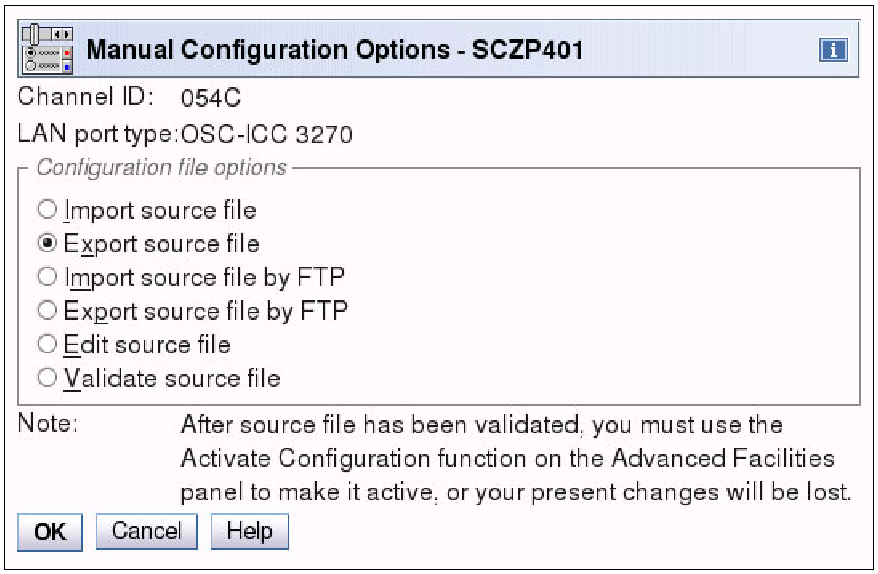

Figure 14-75 OSA Advanced Facilities: export source file

6. The task requests that a file name be written to the installed media device. For our example, we enter `SCZP401_54C` (Figure 14-76). Click **OK**.

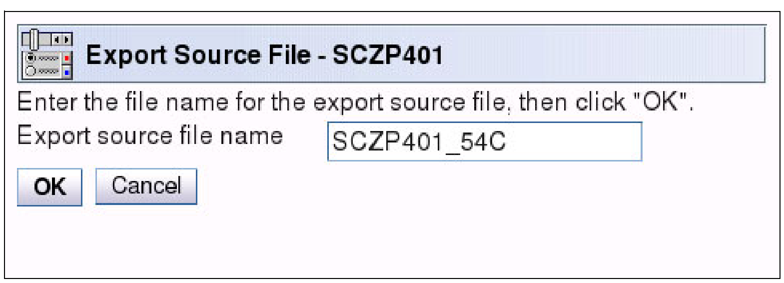

Figure 14-76 OSA Advanced Facilities: export file name specification

7. The HMC displays the ACT20421 window (Figure 14-77). Click **OK**.

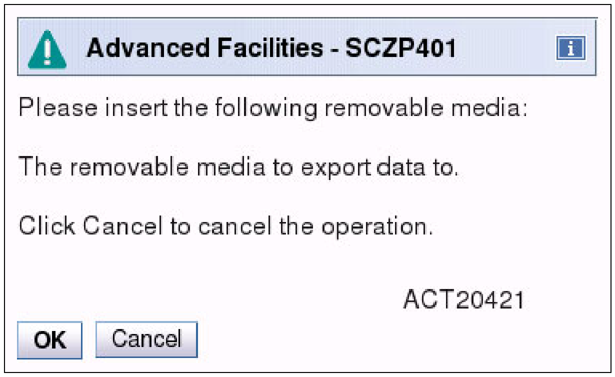

Figure 14-77 ACT20421 window display

8. The HMC task writes the source file for the PCHID that was selected onto the media device and displays a message when it completes (Figure 14-78). Click **OK**.

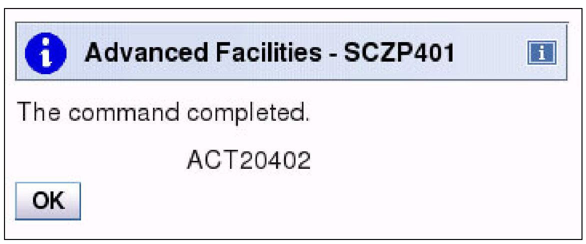

Figure 14-78 OSA Advanced Facilities: export source file completed

9. You can now remove the USB flash memory drive at any time. When you do, the HMC message advises you that the USB flash drive is removed (Figure 14-79). Click **OK**.

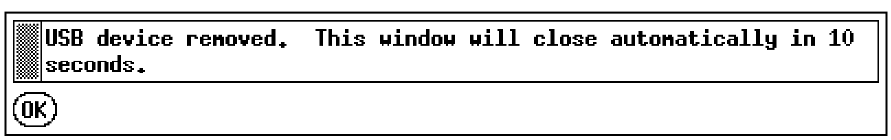

Figure 14-79 OSA Advanced Facilities: removable media removed

10. Click **Cancel** to exit all of the OSA Advanced Facilities windows.

Example 14-2 shows a the source file on the USB flash drive.

Example 14-2 OSA-ICC - sample source file

```
<OSC_SERVER>
<OSC_PHYSICAL_PORT0>
  HOST_IP= 9.12.4.218
  SUBNET_MASK= 255.255.240.0
  PORT= 3270
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= OSAF400
</OSC_PHYSICAL_PORT0>

<OSC_PHYSICAL_PORT1>
  HOST_IP= 192.168.0.1
  SUBNET_MASK= 255.255.255.0
  PORT= 3271
  ETHERNET_FRAME= DIX
  MTU= 1492
  NAME= OSAAF402
</OSC_PHYSICAL_PORT1>

  DEFAULT_GATEWAY= 9.12.4.1
</OSC_SERVER>

<CONFIG_SESSION>
<SESSION1>
  CSS= 00 IID= 02 DEVICE= F400
  GROUP=  "SCC02400"
  CONSOLE_TYPE= 1    RESPONSE= OFF    READ_TIMEOUT= 60
</SESSION1>
...
...
</CONFIG_SESSION>
```

Editing the source file for OSA-ICC

When your OSA-ICC configuration for z13 (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 the following publications, which you can download from IBM Resource Link:

- *OSA-Express Integrated Console Controller User's Guide,* SA22-7990
- *OSA-Express3 Integrated Console Controller Dual-Port User's Guide,* SA23-2266

Importing the OSA-ICC source file to IBM z13 OSC channel

To import the source file and activate the configuration of the OSC-ICC, 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 shown in Figure 14-74 on page 690 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.** This is a same action as shown in step 1 on page 655.
4. Select **Manual configuration options** and click **OK.**

5. The Manual Configuration Options panel opens. Select **Import source file** and click **OK** (Figure 14-80).

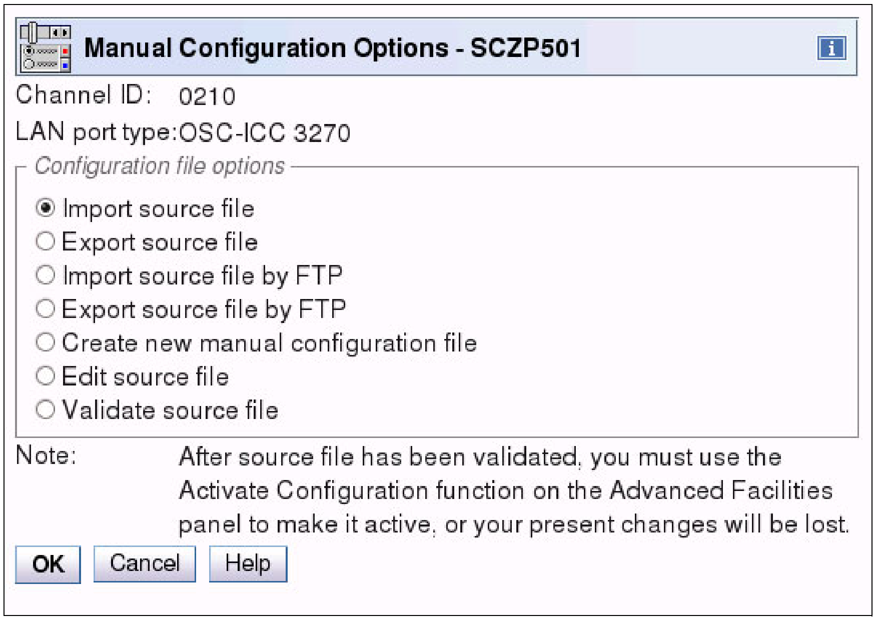

Figure 14-80 OSA Advanced Facilities: export source file

6. You are prompted to insert the media (Figure 14-81). Click **OK**.

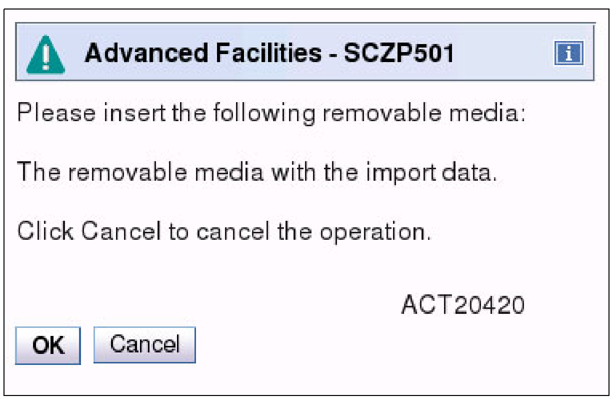

Figure 14-81 OSA Advanced Facilities: insert media

7. The Import Source File window opens. Select the name of the source file you want to import. For the example, we select **SCZP401_54C** as the source file to import (Figure 14-82 on page 695). Click **OK**.

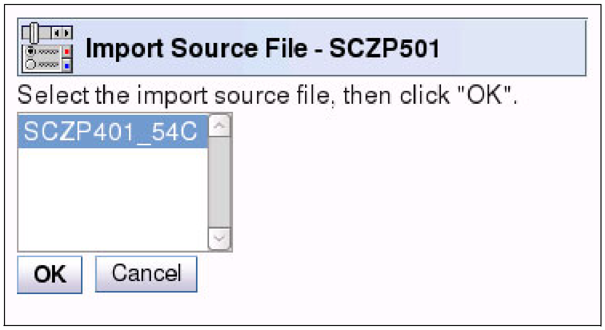

Figure 14-82 OSA Advanced Facilities: import source file selection

8. The next window (Figure 14-83) indicates that the source file import is complete. You can now remove the USB flash drive.

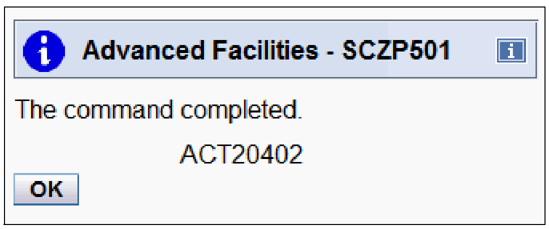

Figure 14-83 OSA Advanced Facilities: Import file success

Validating the source file and activating the configuration

After you import the source file, validate it. If no error was found on the source, you can activate the configuration. If an error was found, you must correct the error.

To validate the imported source file, complete the following steps:

1. In the Manual Configuration Options window, select **Validate source file** and click **OK** (Figure 14-84).

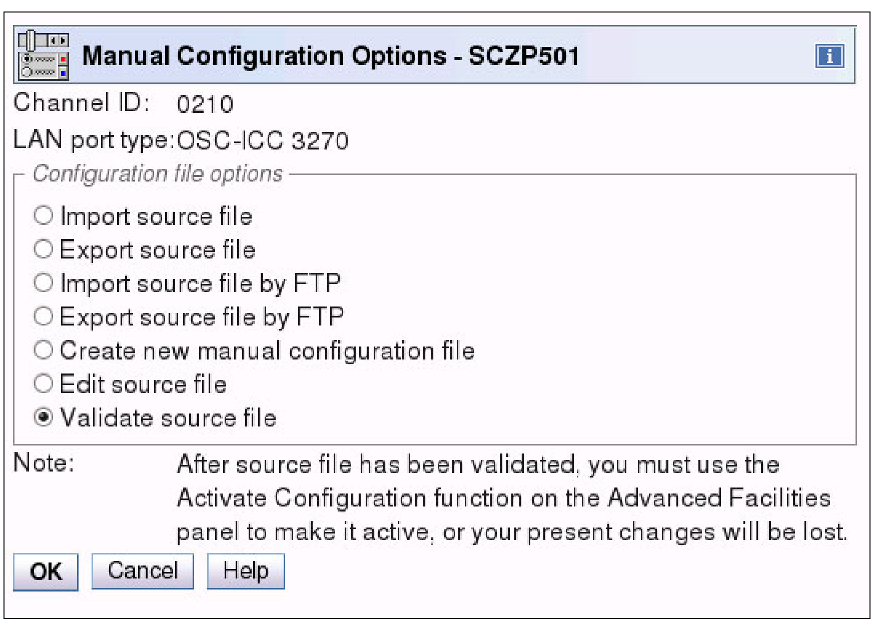

Figure 14-84 OSA Advanced Facilities: validate source file

2. If no error is found, a panel opens (Figure 14-85). Click **OK**, and then click **Cancel.**

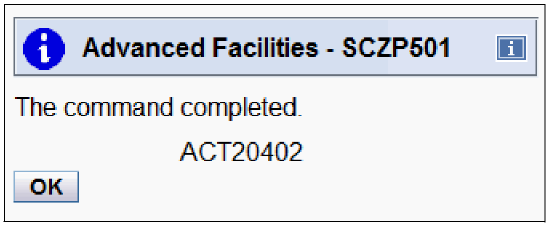

Figure 14-85 OSA Advanced Facilities: source validation success

3. After you validate the source file, activate the configuration by selecting **Activate configuration** in the Advanced Facilities window (Figure 14-86).

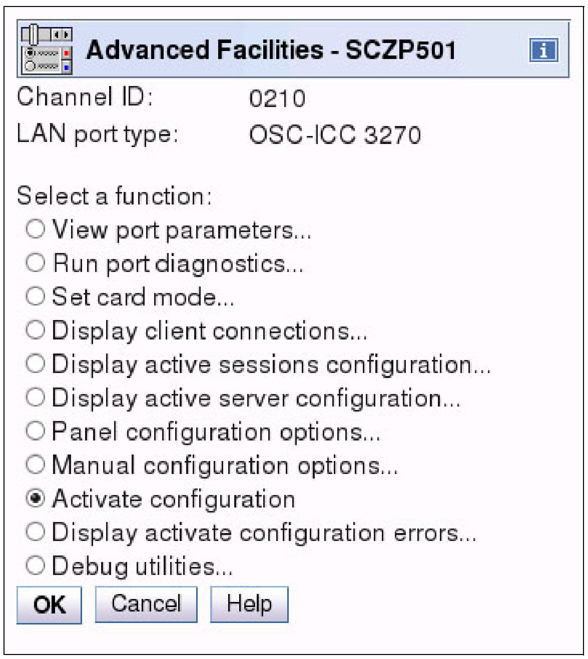

Figure 14-86 OSA Advanced Facilities: activate configuration

4. The confirmation panel opens (Figure 14-87). Click **Yes.**

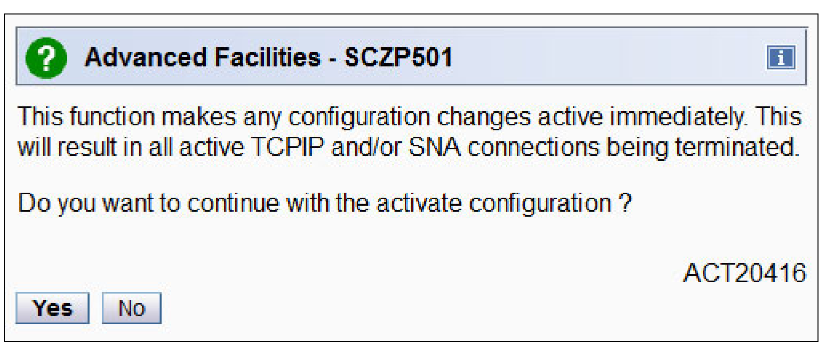

Figure 14-87 OSA Advanced Facilities: confirmation of activate configuration

When OSA-ICC activation succeeds, the panel shown in Figure 14-88 opens.

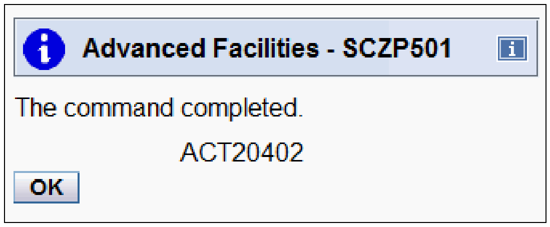

Figure 14-88 OSA Advanced Facilities: activate configuration ended

14.5.3 Define new OSA-ICC configuration

If you newly install z13, you must configure OSA-ICC from the beginning. To define a new OSA-ICC configuration, complete the following steps:

1. Log on to the HMC, select the CPC you want to operate, and open OSA Advanced facilities.
2. Select the OSC PCHID to which you want to define the new OSA-ICC configuration. Then Select **Card specific advanced facilities**. This is the same as step 1 on page 655.
3. Select **Panel configuration options** and click **OK** (Figure 14-89).

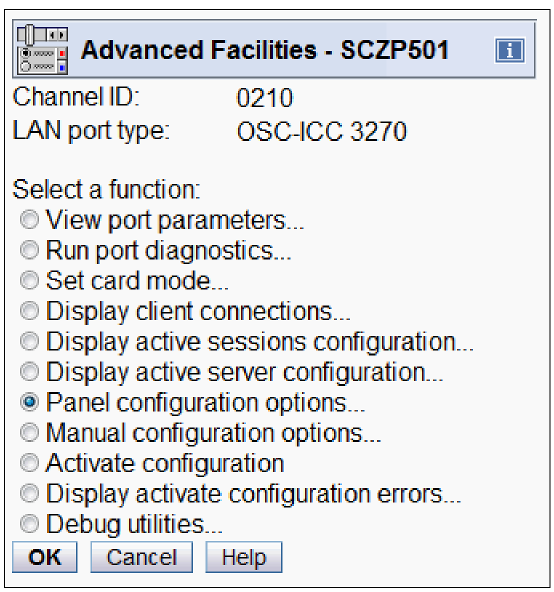

Figure 14-89 OSA Advanced Facilities: Panel configuration options

The Panel Configuration Options window opens (Figure 14-90). This is where you can define session configuration and server configuration. Then, you must validate those values.

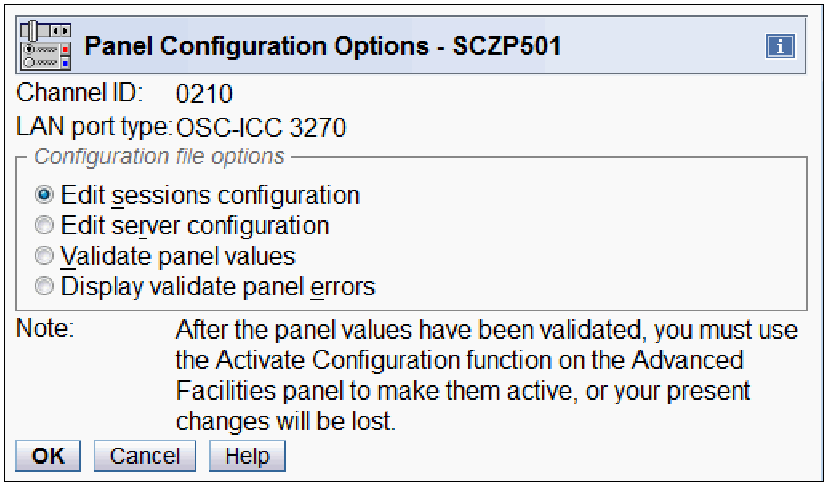

Figure 14-90 OSA Advanced Facilities: Panel Configuration Options

4. In this example, we define the server configuration first. Next, we define the session configuration. To edit server configuration, select **Edit server configuration** and click **OK**.

5. The Edit Server Configuration panel opens. Enter the necessary values on this panel. Figure 14-91 shows our sample configuration. Click **OK** to save.

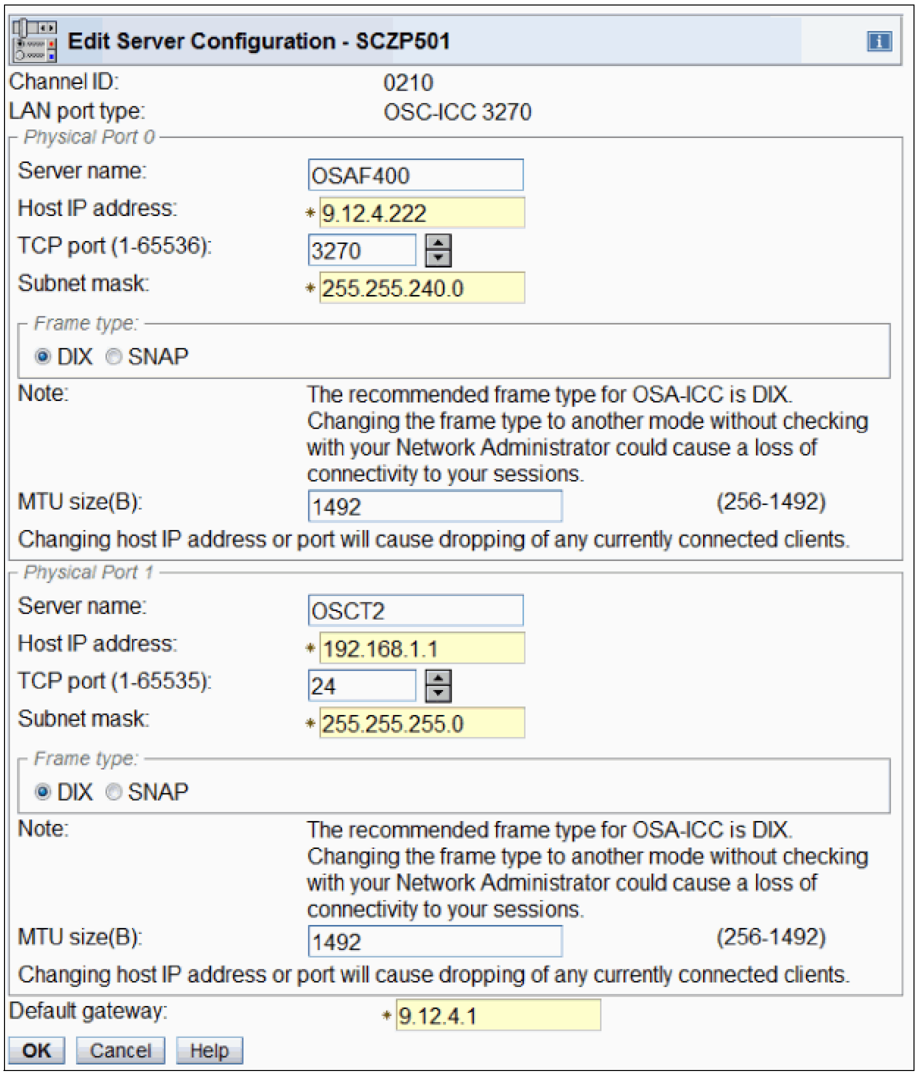

Figure 14-91 Edit Server Configuration

6. The command is completed (ACT20402) window opens. Click **OK** to return to the Panel Configuration Options window (Figure 14-90 on page 698). Next, you define the session configuration.
7. Select **Edit sessions configuration** and click **OK**.

8. The Edit Sessions Configuration window opens (Figure 14-92). To configure a new session, select a number from the Session Index column and click **Change**.

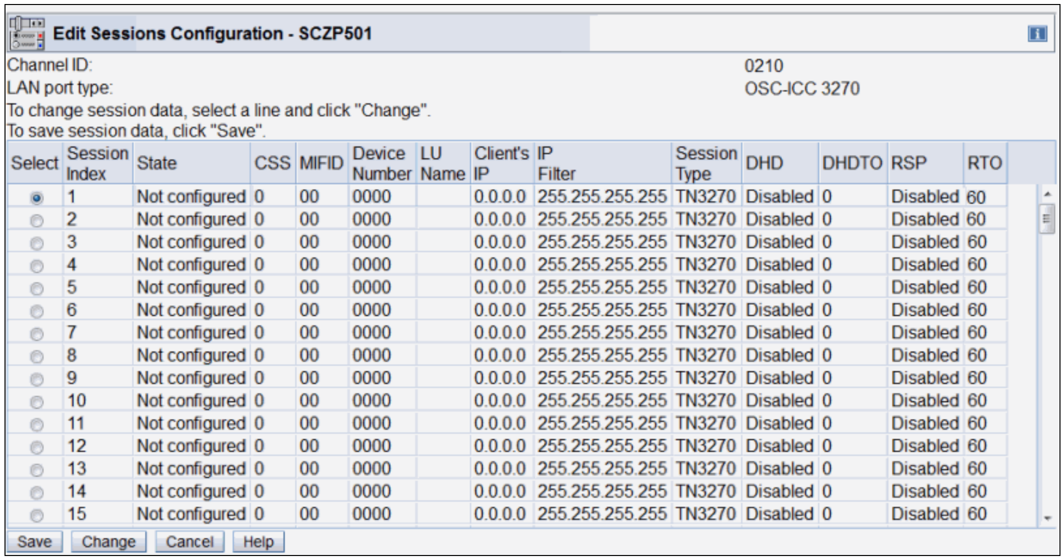

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client's IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
	1	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	2	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	3	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	4	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	5	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	6	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	7	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	8	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	9	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	10	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	11	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	12	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	13	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	14	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
	15	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Figure 14-92 OSA Advanced Facilities: Edit Sessions Configuration

9. The Edit Session Configuration panel opens (Figure 14-93). On this panel, define the value for session here. We use the value shown in the figure. Click **OK** to save.

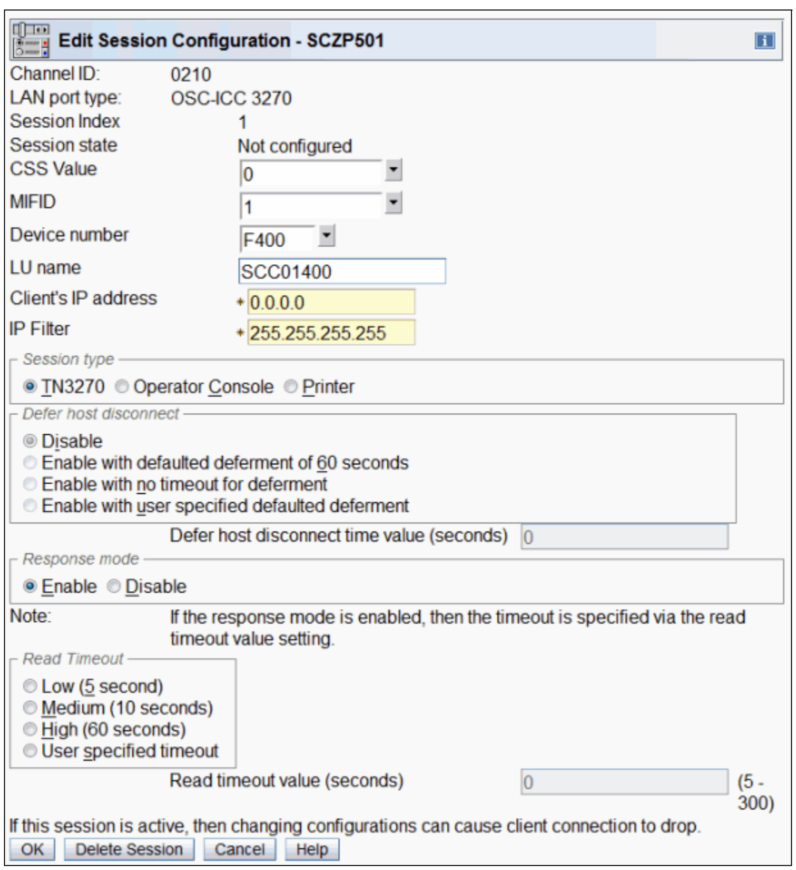

Figure 14-93 OSA Advanced Facilities: Edit Session Configuration

10. You are returned to the Edit Sessions Configuration panel again (Figure 14-94). Be sure that your input values are displayed correctly. To save session values, click **Save**.

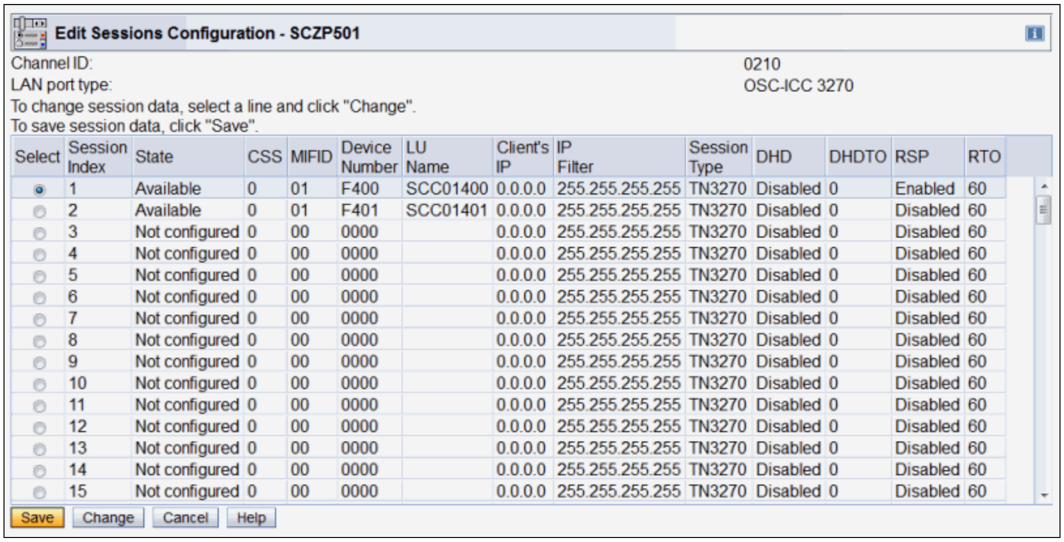

Edit Sessions Configuration - SCZP501

Channel ID: 0210

LAN port type: OSC-ICC 3270

To change session data, select a line and click "Change".

To save session data, click "Save".

Select	Session Index	State	CSS	MIFID	Device Number	LU Name	Client's IP	IP Filter	Session Type	DHD	DHDTO	RSP	RTO
◉	1	Available	0	01	F400	SCC01400	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Enabled	60
○	2	Available	0	01	F401	SCC01401	0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	3	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	4	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	5	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	6	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	7	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	8	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	9	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	10	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	11	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	12	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	13	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	14	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60
○	15	Not configured	0	00	0000		0.0.0.0	255.255.255.255	TN3270	Disabled	0	Disabled	60

Save Change Cancel Help

Figure 14-94 OSA Advanced Facilities: Edit Sessions Configuration after define value

11. The command completed (ACT20402) window opens. Click **OK**.

12. You are returned to the Panel Configuration Options panel (Figure 14-95). Now you can validate. Select **Validate panel values** and click **OK**.

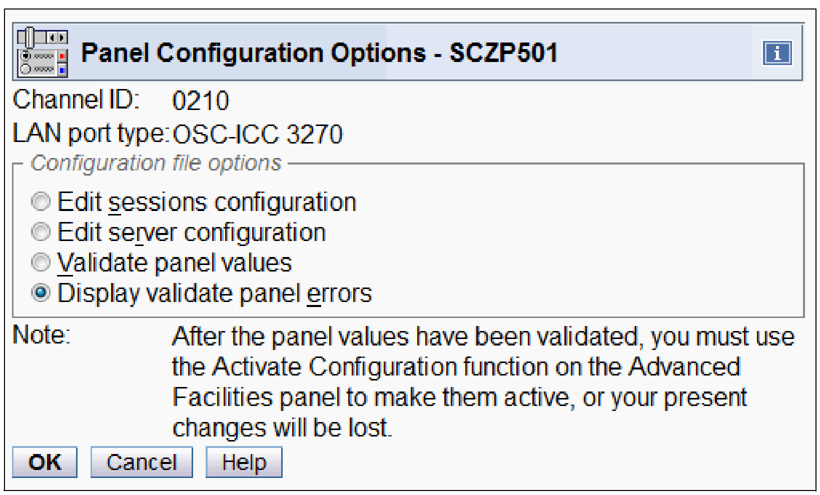

Figure 14-95 OSA Advanced Facilities: Display validate panel errors

If no error is found, the command completed (ACT20402) window opens.

If an error is found, a window similar to Figure 14-96 on page 702 opens. You can confirm the error by selecting **Display validate panel errors** on Panel Configuration Options (Figure 14-95). Correct the error, and again select **Validate panel values** to see that the error is corrected.

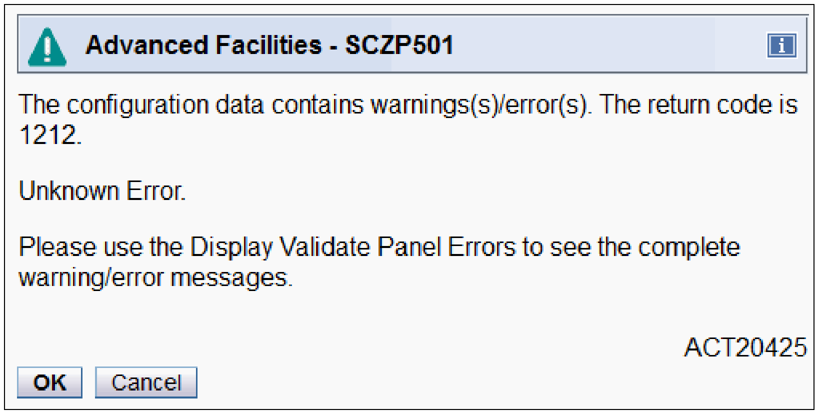

Figure 14-96 OSA Advanced Facilities: Panel Configuration Options error found

13. When you finish validating, you can activate the OSA-ICC configuration. To activate configuration, see step 3 on page 696.

14.6 Setting OSA parameters by using OSA Advanced Facilities

If you want to change the port speed or MAC address of an OSA-Express feature, you can change them through the OSA Advanced Facilities.

Port speed: You can only change port speed to OSA-Express5S 1000BASE-T and OSA-Express4S. Other OSA features have no capability to change port speed.

With OSA-Express5S 1000BASE-T and OSA-Express4S 1000BASE-T, you cannot set port speed to `1000 Mb, Full Duplex` like the OSA- Express2. If you want to set port speed to 1000 Mb, select **Auto Negotiate**.

14.6.1 Setting the OSA port speed by using Advanced Facilities

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 OSD channel for which you need to set the card mode. For the example, we select PCHID 0134 on SCZP501 (Figure 14-97).

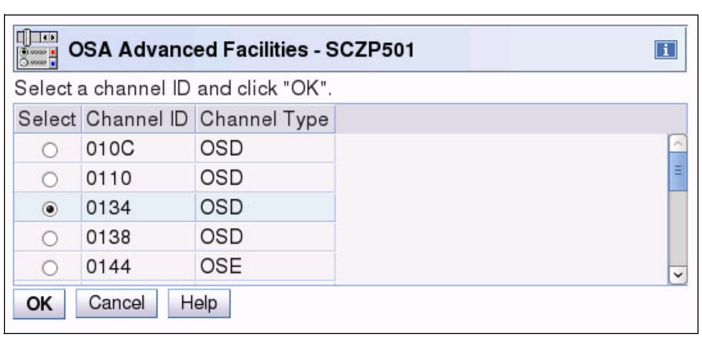

Figure 14-97 OSA Advanced Facilities: PCHID selection

3. Select **Card specific advanced facilities** and click **OK** (Figure 14-98).

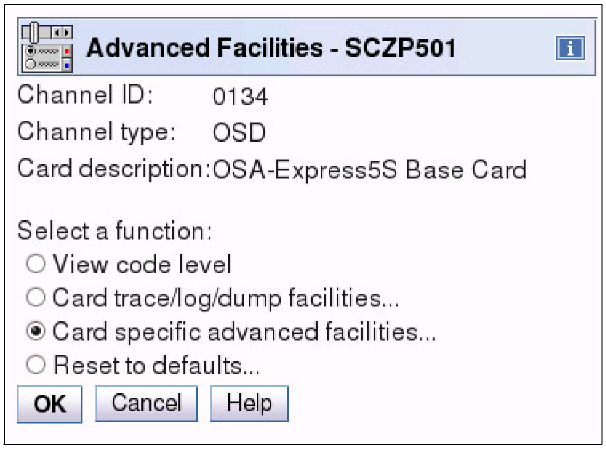

Figure 14-98 OSA Advanced Facilities: card specific advanced facilities

4. Select **Set card mode** and click **OK** (Figure 14-99).

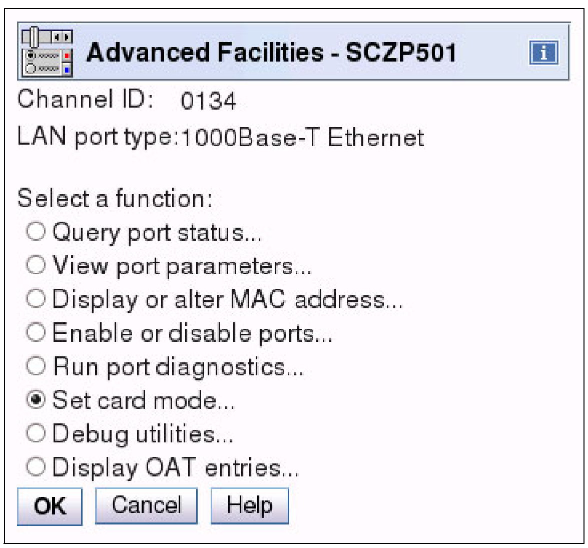

Figure 14-99 Selecting Card Mode definition

5. The Set Card Mode or Speed panel opens (Figure 14-100). Select the correct port speed from the Speed/Mode list. You can set the speed of port 0 and 1 individually. Click **OK**.

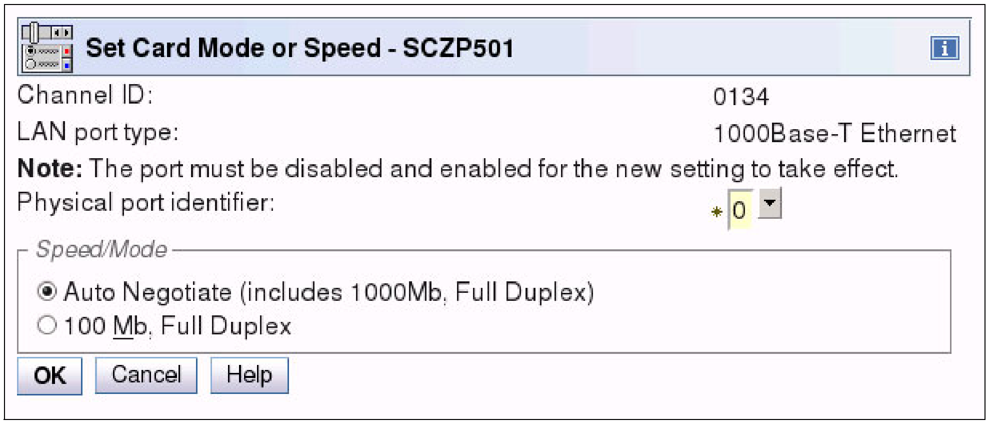

Figure 14-100 OSA Advanced Facilities: set card mode or speed

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

14.6.2 Changing OSA MAC address by using OSA Advanced Facilities

To change the media access control (MAC) address, complete the following steps:

1. Log on to the HMC, then open OSA Advanced Facilities, and select the PCHID you want to customize.
2. The Advanced Facilities panel opens (Figure 14-101). Select **Display or alter MAC address** and then click **OK**.

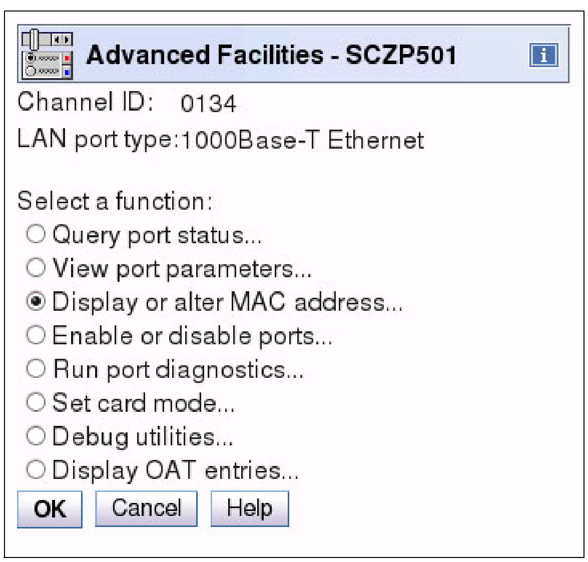

Figure 14-101 OSA Advanced Facilities: display or alter MAC address

3. The Display or alter MAC address window opens (Figure 14-102). Set the MAC address that you want and then click **OK**.

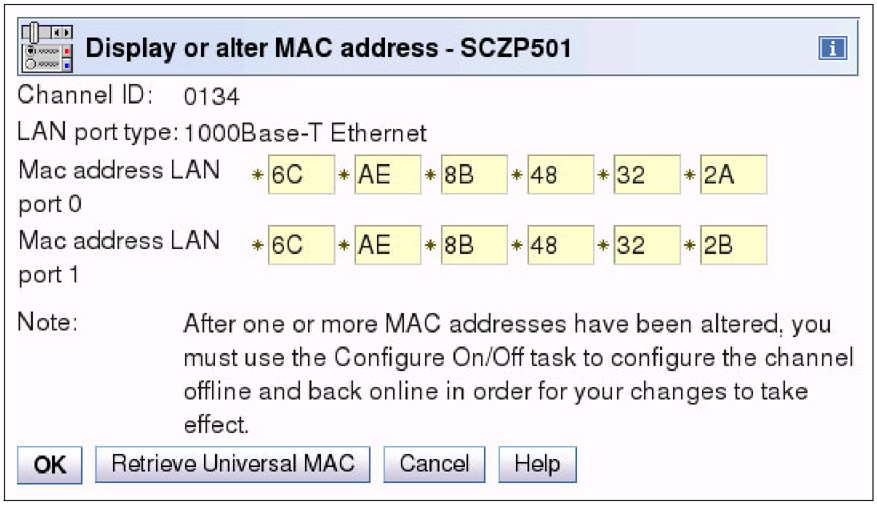

Figure 14-102 OSA Advanced Facility: display or alter MAC address values

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

A

Additional features and functions

This appendix describes the following functions:

- Absolute Capping
- Thin Interrupts
- zEDC
- Simultaneous multithreading (SMT) setup for z/OS

Absolute Capping

Absolute Capping is enabled on z Systems through Licensed Internal Code Version 2.12.1 at Driver Level 15 or later. Absolute Capping is an option that limits the amount of physical processor capacity consumed by an individual LPAR when a PU is defined as a general purpose CP, zIIP, IFL, or an ICF processor that is shared across a set of LPARs.

Originally, Absolute Capping was designed for non z/OS LPARs such as z/VM and Linux on z Systems that do not make use of z/OS Workload Manager (WLM) capacity management functions. Often, Linux customers face the problem that software pricing is based on a “per core” basis, and they need the ability to limit an LPAR to a specific amount of hardware processor capacity in a way that is unaffected by subsequent physical or logical configuration changes. Traditional LPAR capping methodologies do not meet these requirements:

- LPAR defined capacity specified in MSUs (“soft capping”) is supported only for LPARs running z/OS.
- LPAR capping based on weight (“hard capping”) is relative, not absolute. Activating or deactivating other LPARs change the relative weight of all other activated LPARs, and also do dynamic changes to the weight of any running LPAR.

With Absolute Capping, PR/SM establishes a hard capacity limit on the amount of processing capacity that is available to an LPAR. This is done with the specification of an absolute capacity limit for an individual LPAR in absolute processor capacity. The value is actually to hundredths of a processor (for example 4.56 processors) worth of capacity. This setting is specified independently by processor type, namely CPs, zIIPs, IFLs, and also ICFs.

However, the LPAR’s processing weight still dictates the logical partition priority relative to other active partitions sharing the PUs. Therefore Absolute Capping is most effective for caps higher than what the LPAR’s weight would otherwise ensure.

Because Absolute Capping is managed by PR/SM, it is available to all LPARs including z/OS. However, Absolute Capping is ideal for processor types and operating systems that the z/OS WLM cannot control. Absolute Capping is not intended as a replacement for defined capacity or group capacity for z/OS which are managed by WLM.

The specifications for Absolute Capping are done either from the HMC or the SE, by using the Customize/Delete Activation Profile task, which is in the Operational Customization group. Absolute Capping is specified in the image profile of the LPAR in the Processor section, and the next time the LPAR is activated these specifications will be used. This is illustrated in Figure A-1 on page 709, where Absolute Capping for LPAR A06 is enabled and the number of general purpose CPs is set to 1.42. Consider that this value is less than the number of shared CPs that the LPAR is allowed to use, and probably higher than the number of CPs that the weight setting guarantees to the LPAR. These specifications apply only to the CPs used by the LPAR.

For enabling Absolute Capping to the IFLs that LPAR A06 also is using in shared mode, the equivalent specifications must be done by selecting the **Integrated facilities for Linux (IFLs)** check box.

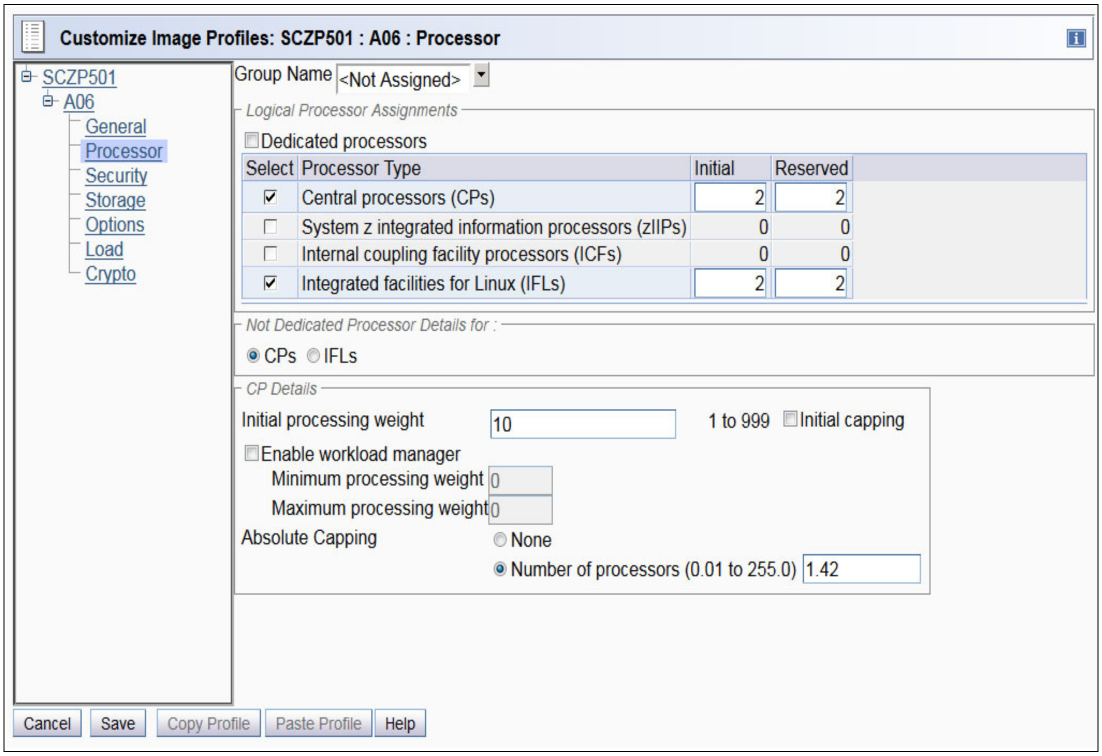

Figure A-1 Absolute Capping Image Profile Specification

If the LPAR is already activated and Absolute Capping must be activated or modified concurrently, this can be done by using Change LPAR Controls task, which also is in the Operational Customization group. Figure A-2 shows the panel for all activated LPARs sharing IFLs. Absolute Capping for IFLs is used in none of the two LPARs A04 and A06. To enable it for LPAR A06, click **None** in the Absolute Capping column.

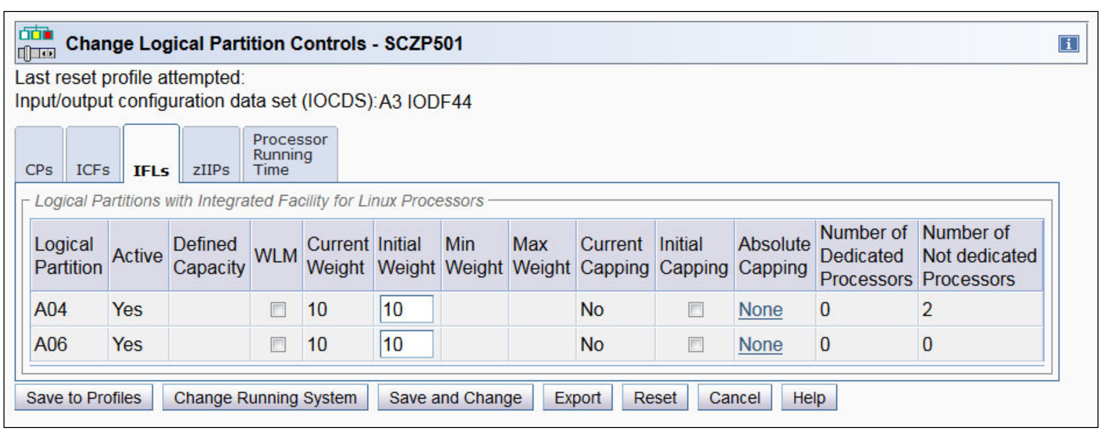

Figure A-2 Change LPAR Controls for Absolute Capping for IFLs - before changes

This leads to the panel shown in Figure A-3, where you can specify a value.

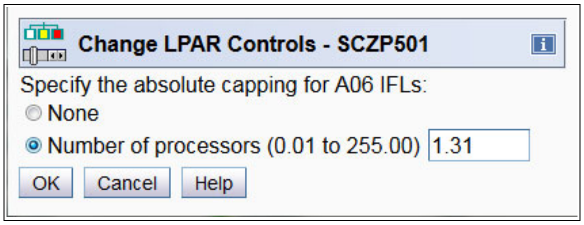

Figure A-3 Change LPAR Controls for Absolute Capping for IFLs - specifying LPAR A06

Click **OK**. The LPAR list is displayed again (Figure A-4), this time showing the newly specified value in the row for LPAR A06.

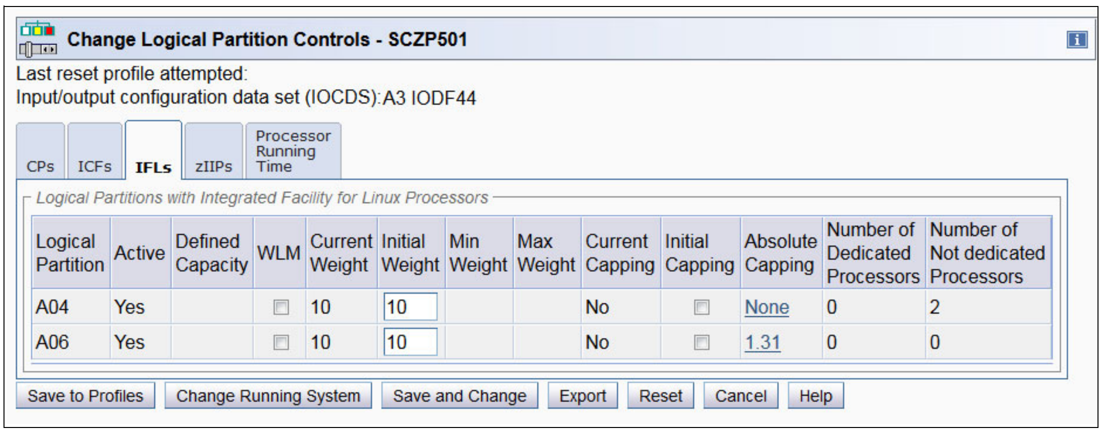

Figure A-4 Change LPAR Controls for Absolute Capping for IFLs - after changes

Absolute Capping is now enabled for the IFLs that are shared by LPAR A06. To enable or change the usage of Absolute Capping for CPs or other processor types, do the same actions for the appropriate processor type. Do not forget to save these changes to the image profiles by clicking either **Save to Profile** or **Save and Change**, if the specifications should also be used after the next activation of the LPAR.

Thin Interrupts

In a Parallel Sysplex environment it is critical for the performance of the applications, that the Coupling Facility is responding as fast as possible to CF requests. Therefore the Coupling Facility Control Code (CFCC), which is the CF operating system implemented in the z Systems firmware, is designed to poll all the time and to consume the available processor resources at 100%. For this reason, be sure to use dedicated CF processors in environments, where the absolute best performance is required.

However, for Parallel Sysplexes with testing or application development workload, a practical approach might be to share CF processors between two or even more CF LPARs. For doing this, modify the polling behavior of the CFCC. This can be done by using the `DYNDISP` command at the CFCC console. The `DYNDISP ON` command prevents the CFCC from polling permanently and causes a dynamic dispatching behavior instead. The `DYNDISP OFF` command leaves the CFCC in the original polling behavior.

CF Level 19 introduced the support for Coupling Thin Interrupts on zBC12 and zEC12 servers, which offers a third option to modify the polling behavior of the CFCC. When using

Thin Interrupts, A CF LPAR is allowed to voluntarily give up control of the shared coupling facility processor whenever it runs out of work to do, and later expediting the redispatching of the partition to allow it to resume the polling logic that is optimal for CFCC to locate and process work. In most shared processor configurations, coupling facilities can achieve a performance benefit when Coupling Thin Interrupts are enabled. This is described in a white paper, which is available at the IBM Techdocs website:

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102400

Thin Interrupts can be turned on and off by issuing the **DYNDISP THIN** command at the CFCC console, which is displayed in the Operating System Messages window for the CF LPAR on the HMC or SE. This is illustrated in Figure A-5 for the CF LPAR A0E, which is running two ICF processors in shared mode.

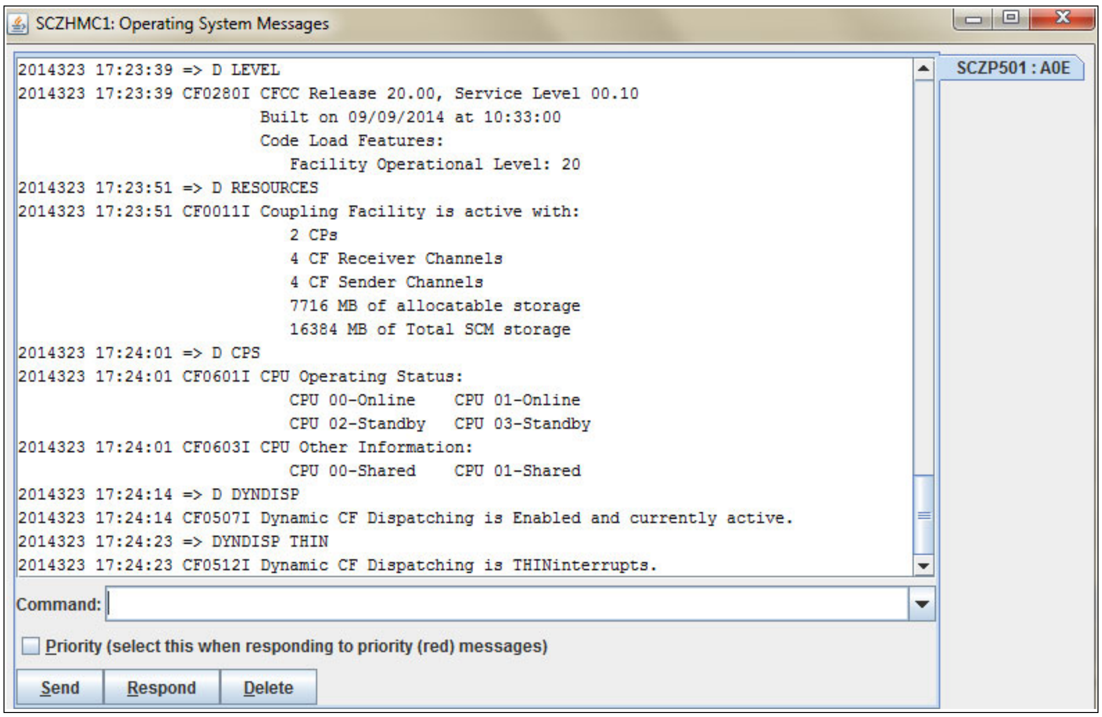

Figure A-5 Operating System Messages for LPAR A0E

zEDC

The IBM zEnterprise Data Compression (zEDC) Express is an optional feature available with the zEC12, zBC12, and z13. 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 two zEDC Express features per drawer domain. Between one and eight features can be installed on the system. There is one PCIe adapter/compression coprocessor per feature, and a zEDC Express feature can be shared by up to 15 LPARs.

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. Additionally, for the full zEDC benefit, zEDC should be active on all systems that might access or share compressed format data sets. This eliminates instances where software inflation would be used when zEDC is not available.

Exploitation support of the zEDC Express feature is provided by z/OS V2R1 for both data compression and decompression. Support for data recovery (decompression), if zEDC is not available on the system, is provided through software on z/OS V2R1, V1R13, and V1R12 with the appropriate PTFs. But because software decompression is slow and uses considerable processor resources, it is not preferred for production environments.

A detailed description about the zEDC Express feature and how to use it is in *Reduce Storage Occupancy and Increase Operations Efficiency with IBM zEnterprise Data Compression*, SG24-8259.

In this appendix, however, only a short summary about the configuration and setup of the zEDC Express feature is given. Several tasks must be done to use and exploit zEDC features:

1. Planning the installation:
   - Consider how many zEDC features are needed.
   - Consider how to share them between the LPARs.
   - Obtain the Physical Channel IDs of the zEDC Express features to be used for HCD definitions.
   - Plan for IPLs to enable the z/OS systems to recognize and use the zEDC features.
2. z/OS: Verifying the prerequisites:
   - Ensure that z/OS 2.1 together with the latest fixes is installed.
   - Ensure that on all z/OS 1.13 or 1.12 the required PTFs for data recovery (decompression) are installed.
3. z/OS: Enabling the Priced Software Feature:
   - For exploitation of the zEDC Express feature, a Priced Software Feature for z/OS must be enabled.
   - Update the IFAPRDxx PARMLIB member in z/OS 2.1.
   - IPL the z/OS system so that the zEDC Express feature's device driver can recognize the enablement.
4. z/OS: Controlling the use of PCIe features:
   - Update the IQPPRMxx parmlib member in z/OS 2.1.
5. HCD: Defining the PCIe feature:
   - Define the zEDC device in HCD. There are a few differences in defining PCIe features compared to the usual I/O devices.
6. HCD: Activating the new configuration:
   - The typical procedure is to build a production IODF and activate the new configuration dynamically.
7. z/OS: Bringing the zEDC Express PCIe features online:
   - Make sure that z/OS has access to the newly defined PCIe features.
   - Configure them online.
8. z/OS: Managing the zEDC Express PCIe features:
   - Use the usual commands and procedures to display, configure, and toggle the zEDC feature.

Simultaneous multithreading (SMT) setup for z/OS

This section describes how to set up simultaneous multithreading (SMT) in z/OS V2.1.

z13 supports two-way simultaneous multithreading (SMT2) operation. z/OS can support SMT2 for zIIPs. Up to two active execution threads per core can dynamically share the caches, TLBs, and execution resources of each zIIP core. SMT is designed to improve both core capacity and single thread performance significantly.

When the SMT facility is enabled, the hardware threads within a core share certain hardware resources such as execution units and caches. When the SMT is not enabled, a core executes a single hardware thread.

Prerequisite APAR for SMT2

To use the SMT2 function on z/OS, these APARs must be applied:

- OA43366 (BCP)
- OA43622 (WLM)
- OA44439 (XCF)

Define PROCessor VIEW (PROCVIEW) in LOADxx

To activate SMT2 function on z/OS, define the **PROCVIEW CORE** option in LOADxx. If you do not use the SMT2 function, define **PROCVIEW CPU**. The default parameter of **PROCVIEW** is **CPU**.

When you define **PROCVIEW CORE**, you cannot use the word **CPU** on the z/OS command. You must use **CORE** instead of **CPU**. If you want to continue to use **CPU**, define **PROCVIEW CORE,CPU_OK**. This parameter causes z/OS to treat CPU as an acceptable alias for CORE. Example A-1 is our LOADxx definition to activate SMT2 function.

> **Note:** Fall back from PROVIEW CORE to PROCVIEW CPU requires IPL.

Example: A-1 LOADxx definition sample with CORE,CPU_OK

```
HWNAME   SCZP501
LPARNAME A01
SYSPLEX  PLEX76   Y
IODF     ** SYS6     L06RMVS1 01 Y
SYSCAT   BH6CAT123CMCAT.BH6CAT
PARMLIB  SYS1.PARMLIB
PARMLIB  SYS1.IBM.PARMLIB
PROCVIEW CORE,CPU_OK
```

Define MT_ZIIP_MODE=2 in IEAOPTxx

To execute multiple threads on zIIP SMT mode, define the **MT_ZIIP_MODE=2** parameter.

MT_ZIIP_MODE=2 means for two active threads (the default is 1).

- When **PROCVIEW CPU** is specified, the processor class MT Mode is always 1.
- Without an IPL you can change the zIIP processor class MT Mode (the number of active threads per online zIIP) using **IEAOPTxx. SET OPT=xx**.
- When you define **MT_ZIIP_MODE=2**, Hiperdispatch is automatically set YES.

Example A-2 is our IEAOPTxx sample to activate SMT2 on zAAP.

Example: A-2 IEAOPTxx definition with MT_ZIIP_MODE=2

```
CPENABLE=(10,30)
HIPERDISPATCH=YES
MT_ZIIP_MODE=2
```

CPU Address changes with SMT2

When you activate SMT2 function, the CPU address changes. A 16-bit CPU ID consists of a 15-bit Core ID and 1-bit Thread ID (Figure A-6).

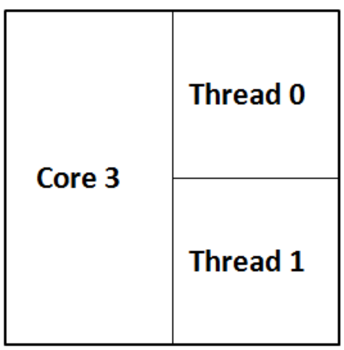

Figure A-6 z13 processor chip: Core Virtualization

Consider this information:

- CPU ID 6 (b'000000000000011**0**') means core 3 Thread **0**.
- CPU ID 7 (b'000000000000011**1**') means core 3 Thread **1**.

z/OS D M=CORE, D M=CPU command display

In this section, we show various **D M=CORE** and **D M=CPU** command responses.

D M=CPU command with PROCVIEW CPU

Figure A-7 is a command response of the **D M=CPU** command. In this example, we define PROCVIEW CPU.

```
D M=CPU
IEE174I 15.06.19 DISPLAY M 545
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DA872964
01  +                     01DA872964
02  +                     01DA872964
03  +                     01DA872964
04  +I                    01DA872964
05  -
06  -
07  -
08  -
09  -I

CPC ND = 002964.N63.IBM.02.00000008DA87
CPC SI = 2964.735.IBM.02.000000000008DA87
Model: N63
CPC ID = 00
CPC NAME = SCZP501
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE

I        INTEGRATED INFORMATION PROCESSOR (zIIP)
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER
```

Figure A-7 D M=CPU command with PROCVIEW CPU

D M=CORE command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=2

Figure A-8 is a command response of the **D M=CORE** command. In this example, we define PROCVIEW CORE,CPU_OK in LOADxx, and MT_ZIIP_MODE=2 in IEAOPTxx.

```
D M=CORE
IEE174I 14.54.21 DISPLAY M 902
CORE STATUS: HD=Y   MT=2  MT_MODE: CP=1  zIIP=2
ID    ST   ID RANGE   VP  ISCM  CPU THREAD STATUS
0000   +   0000-0001  H   0000  +N
0001   +   0002-0003  H   0000  +N
0002   +   0004-0005  H   0000  +N
0003   +   0006-0007  H   FC00  +N
0004   +I  0008-0009  H   0200  ++
0005   -   000A-000B
0006   -   000C-000D
0007   -   000E-000F
0008   -   0010-0011
0009   -I  0012-0013

CPC ND = 002964.N63.IBM.02.00000008DA87
CPC SI = 2964.735.IBM.02.000000000008DA87
         Model: N63
CPC ID = 00
CPC NAME = SCZP501
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    N NOT AVAILABLE    / MIXED STATE
W WLM-MANAGED

I        INTEGRATED INFORMATION PROCESSOR (zIIP)
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER
```

Figure A-8 .D M=CORE command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=2

D M=CPU command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=2

Figure A-9 is a command response of the **D M=CPU** command. In this example, we define PROCVIEW CORE,CPU_OK in LOADxx, and MT_ZIIP_MODE=2 in IEAOPTxx. You can confirm that word CPU is valid in the command.

```
D M=CPU
IEE174I 14.56.02 DISPLAY M 904
CORE STATUS: HD=Y   MT=2  MT_MODE: CP=1  zIIP=2
ID    ST   ID RANGE   VP  ISCM  CPU THREAD STATUS
0000   +   0000-0001  H   0000  +N
0001   +   0002-0003  H   0000  +N
0002   +   0004-0005  H   0000  +N
0003   +   0006-0007  H   FC00  +N
0004   +I  0008-0009  H   0200  ++
0005   -   000A-000B
0006   -   000C-000D
0007   -   000E-000F
0008   -   0010-0011
0009   -I  0012-0013

CPC ND = 002964.N63.IBM.02.00000008DA87
CPC SI = 2964.735.IBM.02.000000000008DA87
         Model: N63
CPC ID = 00
CPC NAME = SCZP501
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    N NOT AVAILABLE    / MIXED STATE
W WLM-MANAGED
I        INTEGRATED INFORMATION PROCESSOR (zIIP)
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER
```

Figure A-9 D M=CPU command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=2

D M=CORE command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=1

Figure A-10 is a command response of the `D M=CORE` command. In this example, we define PROCVIEW CORE,CPU_OK in LOADxx, and MT_ZIIP_MODE=2 in IEAOPTxx. You can confirm that status of zIIP changed from “++” to “+N”, which indicates the SMT2 function will not work on this zIIP.

```
D M=CORE
IEE174I 16.04.06 DISPLAY M 209
CORE STATUS: HD=Y   MT=2  MT_MODE: CP=1  zIIP=1
ID    ST   ID RANGE   VP  ISCM  CPU THREAD STATUS
0000   +   0000-0001  H   0000  +N
0001   +   0002-0003  H   0000  +N
0002   +   0004-0005  H   0000  +N
0003   +   0006-0007  H   FC00  +N
0004   +I  0008-0009  H   0200  +N
0005   -   000A-000B
0006   -   000C-000D
0007   -   000E-000F
0008   -   0010-0011
0009   -I  0012-0013

CPC ND = 002964.N63.IBM.02.00000008DA87
CPC SI = 2964.735.IBM.02.000000000008DA87
         Model: N63
CPC ID = 00
CPC NAME = SCZP501
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    N NOT AVAILABLE    / MIXED STATE
W WLM-MANAGED

I        INTEGRATED INFORMATION PROCESSOR (zIIP)
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC SI  SYSTEM INFORMATION FROM STSI INSTRUCTION
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
CPC NAME CENTRAL PROCESSING COMPLEX NAME
LP NAME  LOGICAL PARTITION NAME
LP ID    LOGICAL PARTITION IDENTIFIER
CSS ID   CHANNEL SUBSYSTEM IDENTIFIER
MIF ID   MULTIPLE IMAGE FACILITY IMAGE IDENTIFIER
```

Figure A-10 D M=CORE command with PROCVIEW CORE,CPU_OK and MT_ZIIP_MODE=1

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Extending z/OS System Management Functions with IBM zAware*, SG24-8070
- *IBM z Systems Connectivity Handbook*, SG24-5444
- *IBM z13 Technical Introduction*, SG24-8250
- *IBM z13 Technical Guide*, SG24-8251
- *Implementing and Managing InfiniBand Coupling Links on IBM System z,* SG24-7539
- *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

Other publications

These publications are also relevant as further information sources:

- *IBM System z Advanced Workload Analysis Reporter (IBM zAware) Guide*, SC27-2623
- *Installation Manual for Physical Planning 2458-003*, GC27-2619
- *Installation Manual - Physical Planning 2827,* GC28-6914
- *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152
- *CHPID Mapping Tool User's Guide*, GC28-6947-00
- *z/OS MVS Setting Up a Sysplex,* SA23-1399
- *Open Systems Adapter-Express Customer's Guide and Reference,* SA22-7935
- *Open Systems Adapter/Support Facility on the Hardware Management Console,* SC14-7580
- *OSA-Express Integrated Console Controller User's Guide,* SA22-7990
- *OSA-Express3 Integrated Console Controller Dual-Port User's Guide,* SA23-2266
- *z/OS V2R1.0 Communications Server: IP Configuration Guide,* SC27-3650

- *Hardware Configuration Definition Planning,* GA32-0907
- *Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037
- *System z Application Programming Interfaces*, SB10-7030
- *Hardware Management Console Web Service API,* SB27-2616
- *Common Information Model (CIM) Management Interfaces,* SB10-7154
- *zEnterprise System Capacity on Demand User's Guide,* SC28-2605
- *z/OS MVS Capacity Provisioning User's Guide,* SC33-8299

Online resources

These websites are also relevant as further information sources:

- IBM z Systems

  http://www.ibm.com/systems/z/
- IBM Resource Link:

  http://www.ibm.com/servers/resourcelink
- DMTF:

  http://www.dmtf.org
- PRIZM Protocol Converter Appliance, see Optica Technologies:

  http://www.opticatech.com/
- InfiniBand:

  http://www.infinibandta.org
- CFSizer tool:

  http://www.ibm.com/systems/z/cfsizer
- System Storage Interoperation Center (SSIC):

  http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss
- IBM Techdocs:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Redbooks

IBM z13 Configuration Setup

SG24-8260-00

ISBN 0738440639

IBM

(1.0" spine)
0.875"<->1.498"
460 <-> 788 pages

SG24-8260-00

ISBN 0738440639

Printed in U.S.A.