


# FICON Planning and Implementation Guide

Learn about topologies, concepts, and terminology

Receive guidance for planning, implementation, and migration

Get realistic examples and scenarios



Karan Singh  
Marimar Grau  
Peter Hoyle  
Jonghak Kim  
Jukka Myyryläinen  
Mark Pope  
Jean Sylvester  
Francisco Vilaca

**Red**books





International Technical Support Organization

**FICON Planning and Implementation Guide**

March 2012

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

#### **Fourth Edition (March 2012)**

This edition applies to FICON features defined as CHPID type FC, supporting native FICON, High Performance FICON for System z (zHPF), and FICON Channel-to-Channel (CTC) on IBM zEnterprise 196, zEnterprise 114 and System z10 servers.

**© Copyright International Business Machines Corporation 2012. All rights reserved.**

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

|  |      |
|--|------|
| <b>Notices</b> .....   | ix   |
| Trademarks .....   | x    |
| <b>Preface</b> .....   | xi   |
| The team who wrote this book .....                           | xi   |
| Now you can become a published author, too! .....            | xiii |
| Comments welcome .....                                       | xiv  |
| Stay connected to IBM Redbooks .....                         | xiv  |
| <b>Part 1. Understanding FICON</b> .....                     | 1    |
| <b>Chapter 1. Introduction to FICON</b> .....                | 3    |
| 1.1 Basic Fibre Channel terminology .....                    | 4    |
| 1.1.1 Node .....   | 6    |
| 1.1.2 Port .....   | 6    |
| 1.1.3 Switched fabric .....                                  | 6    |
| 1.1.4 FC link .....  | 7    |
| 1.1.5 World Wide Name .....                                  | 7    |
| 1.1.6 Open exchanges .....                                   | 8    |
| 1.2 System z FICON .....                                     | 9    |
| 1.2.1 High Performance FICON for System z (zHPF) .....       | 10   |
| 1.2.2 FICON-to-ESCON solutions .....                         | 13   |
| 1.2.3 FICON operating modes and topologies .....             | 14   |
| 1.2.4 Platform and name registration in FICON channel .....  | 15   |
| 1.2.5 Terms used with FICON Directors .....                  | 16   |
| 1.2.6 Terms used with the Input/Output architecture .....    | 16   |
| 1.3 z/OS discovery and auto-configuration (zDAC) .....       | 18   |
| <b>Chapter 2. System z FICON technical description</b> ..... | 21   |
| 2.1 Using the FICON architecture for I/O operations .....    | 22   |
| 2.1.1 FICON initialization .....                             | 22   |
| 2.1.2 FICON I/O request .....                                | 24   |
| 2.1.3 Command mode .....                                     | 28   |
| 2.1.4 Transport mode .....                                   | 28   |
| 2.1.5 <b>Missing Interrupt Handler</b> .....                 | 31   |
| 2.1.6 zHPF Temporary Logout .....                            | 32   |
| 2.1.7 Modified Indirect Data Address Word .....              | 32   |
| 2.1.8 Transport Indirect Data Address Word .....             | 33   |
| 2.1.9 Open exchange .....                                    | 34   |
| 2.1.10 Buffer-to-buffer credit usage in FICON .....          | 34   |
| 2.1.11 Extended distance FICON .....                         | 35   |
| 2.1.12 FICON link incident reporting .....                   | 36   |
| 2.2 System z FICON feature support .....                     | 37   |
| <b>Chapter 3. FICON Director technical description</b> ..... | 45   |
| 3.1 The role of the FICON Director .....                     | 46   |
| 3.1.1 Switched configuration .....                           | 46   |
| 3.1.2 Cascaded configuration .....                           | 47   |
| 3.1.3 Basic components of a FICON Director .....             | 48   |

|  |            |
|--|------------|
| 3.1.4 Basic functions of a FICON Director . . . . .                        | 50         |
| 3.2 Qualified FICON Directors . . . . .                                    | 56         |
| 3.2.1 IBM System Storage SAN b-type family components . . . . .            | 57         |
| 3.2.2 IBM System Storage SAN b-type family functions . . . . .             | 63         |
| 3.2.3 Cisco MDS 9500 Series . . . . .                                      | 67         |
| 3.2.4 Functions of the Cisco MDS 9500 Series . . . . .                     | 71         |
| <b>Part 2. Planning the FICON environment . . . . .</b>                    | <b>75</b>  |
| <b>Chapter 4. Planning the FICON environment . . . . .</b>                 | <b>77</b>  |
| 4.1 Structured approach for planning . . . . .                             | 78         |
| 4.2 Documentation . . . . .  | 80         |
| 4.3 Requirements . . . . .   | 81         |
| 4.4 Context . . . . .  | 82         |
| 4.4.1 Migrating from ESCON to FICON . . . . .                              | 82         |
| 4.4.2 Moving to a high bandwidth environment (FICON Express8/8S) . . . . . | 83         |
| 4.4.3 Migrating from a single site to a multi-site environment . . . . .   | 84         |
| 4.4.4 Implementing a new FICON environment . . . . .                       | 85         |
| 4.5 Topologies and supported distances . . . . .                           | 85         |
| 4.5.1 Point-to-point . . . . .   | 86         |
| 4.5.2 Switched . . . . .   | 86         |
| 4.5.3 Cascaded . . . . .   | 87         |
| 4.5.4 Extended distance . . . . .  | 88         |
| 4.6 Convergence . . . . .  | 90         |
| 4.6.1 Intermix fabric . . . . .  | 90         |
| 4.6.2 Fabric security . . . . .  | 91         |
| 4.6.3 High integrity . . . . .   | 92         |
| 4.6.4 Zoning . . . . .   | 93         |
| 4.7 Management . . . . .   | 94         |
| 4.7.1 Command-line interface . . . . .                                     | 94         |
| 4.7.2 Element management . . . . .   | 95         |
| 4.7.3 Fabric management . . . . .  | 95         |
| 4.7.4 Storage management initiative specification . . . . .                | 95         |
| 4.7.5 System z management for FICON Directors . . . . .                    | 96         |
| 4.8 Virtualization and availability . . . . .                              | 97         |
| 4.8.1 System z . . . . .   | 97         |
| 4.8.2 Control unit . . . . .   | 100        |
| 4.8.3 FICON Director . . . . .   | 101        |
| 4.9 Performance . . . . .  | 102        |
| 4.9.1 Frame pacing . . . . .   | 103        |
| 4.9.2 Extended distance FICON . . . . .                                    | 104        |
| 4.9.3 Multiple allegiance . . . . .  | 105        |
| 4.9.4 Parallel Access Volume and HyperPAV . . . . .                        | 105        |
| 4.9.5 Modified Indirect Data Address Word . . . . .                        | 106        |
| 4.9.6 High Performance FICON . . . . .                                     | 107        |
| 4.9.7 Bandwidth management . . . . .                                       | 108        |
| 4.9.8 Traffic management . . . . .   | 111        |
| 4.9.9 Evaluation tools . . . . .   | 111        |
| 4.10 Prerequisites and interoperability . . . . .                          | 112        |
| 4.11 Physical connectivity . . . . .                                       | 113        |
| <b>Part 3. Configuring the FICON environment . . . . .</b>                 | <b>117</b> |
| <b>Chapter 5. Configuring a point-to-point topology . . . . .</b>          | <b>119</b> |

|   |   |            |
|---|---|------------|
| 5.1   | Establishing a point-to-point topology . . . . .                | 120        |
| 5.2   | Description of our environment . . . . .                        | 120        |
| 5.3   | Tasks and checklist . . . . .                                   | 121        |
| 5.3.1   | Tasks . . . . .   | 121        |
| 5.3.2   | Verification checklist . . . . .                                | 121        |
| 5.4   | Defining the new configuration . . . . .                        | 122        |
| 5.4.1   | Defining the channel, CU, and storage devices . . . . .         | 123        |
| 5.4.2   | Defining our configuration example using HCD . . . . .          | 123        |
| 5.4.3   | Defining our configuration example using HCM . . . . .          | 128        |
| 5.4.4   | How the definition looks as IOCP statements . . . . .           | 134        |
| 5.4.5   | Defining a FICON Channel-to-Channel (FCTC) connection . . . . . | 135        |
| 5.5   | Building and activating the IODF . . . . .                      | 137        |
| 5.5.1   | Building the production IODF . . . . .                          | 138        |
| 5.5.2   | Steps to activate the production IODF . . . . .                 | 141        |
| 5.6   | Configuring the IBM Storage System DS8000 . . . . .             | 146        |
| 5.6.1   | Connecting the fiber optic cables . . . . .                     | 147        |
| 5.6.2   | Configuring the channel paths and devices online . . . . .      | 147        |
| 5.6.3   | Verifying the installation . . . . .                            | 149        |
| <b>Chapter 6. Configuring a switched topology . . . . .</b> |   | <b>155</b> |
| 6.1   | Establishing a switched topology . . . . .                      | 156        |
| 6.2   | Description of our environment . . . . .                        | 156        |
| 6.3   | Tasks and checklist . . . . .                                   | 157        |
| 6.3.1   | Tasks . . . . .   | 157        |
| 6.3.2   | Verification checklist . . . . .                                | 158        |
| 6.4   | Defining the new configuration . . . . .                        | 159        |
| 6.4.1   | Defining the channel, CU, and storage devices . . . . .         | 159        |
| 6.4.2   | Defining our configuration example using HCD . . . . .          | 160        |
| 6.4.3   | Defining our configuration example using HCM . . . . .          | 166        |
| 6.4.4   | How the definition looks as IOCP statements . . . . .           | 175        |
| 6.4.5   | Defining FICON Channel-to-Channel (FCTC) connection . . . . .   | 176        |
| 6.5   | Building and activating the IODF . . . . .                      | 179        |
| 6.5.1   | Building the production IODF . . . . .                          | 179        |
| 6.5.2   | Activating the production IODF . . . . .                        | 182        |
| 6.6   | Configuring the IBM Storage System DS8000 . . . . .             | 188        |
| 6.6.1   | Configuring the FICON Director . . . . .                        | 188        |
| 6.6.2   | Connecting the fiber optic cables . . . . .                     | 189        |
| 6.6.3   | Configuring the channel path and device online . . . . .        | 189        |
| 6.6.4   | Verifying the installation . . . . .                            | 192        |
| <b>Chapter 7. Configuring a cascaded topology . . . . .</b> |   | <b>197</b> |
| 7.1   | Establishing a cascaded topology . . . . .                      | 198        |
| 7.2   | Description of our environment . . . . .                        | 198        |
| 7.3   | Tasks and checklist . . . . .                                   | 199        |
| 7.3.1   | Tasks . . . . .   | 199        |
| 7.3.2   | Verification checklist . . . . .                                | 200        |
| 7.4   | Defining the new configuration . . . . .                        | 201        |
| 7.4.1   | Defining the channel, CU, and storage devices . . . . .         | 201        |
| 7.4.2   | Defining our configuration example using HCD . . . . .          | 202        |
| 7.4.3   | Defining our configuration example using HCM . . . . .          | 208        |
| 7.4.4   | How the definition looks as IOCP statements . . . . .           | 218        |
| 7.4.5   | Defining a FICON Channel-to-Channel (FCTC) connection . . . . . | 220        |
| 7.5   | Building and activating the IODF . . . . .                      | 227        |

|                   |   |            |
|-------------------|---|------------|
| 7.5.1             | Building the production IODF . . . . .                        | 227        |
| 7.5.2             | Activating the production IODF . . . . .                      | 230        |
| 7.6               | Configuring the IBM Storage System DS8000 . . . . .           | 236        |
| 7.6.1             | Configuring the FICON Directors . . . . .                     | 236        |
| 7.6.2             | Connecting the fiber optic cables . . . . .                   | 237        |
| 7.6.3             | Configuring the channel path and device online . . . . .      | 237        |
| 7.6.4             | Verifying the installation . . . . .                          | 240        |
| <b>Chapter 8.</b> | <b>Configuring FICON Directors . . . . .</b>                  | <b>245</b> |
| 8.1               | Configuration overview . . . . .                              | 246        |
| 8.1.1             | Configuration flowchart . . . . .                             | 246        |
| 8.1.2             | FICON Director management connectivity . . . . .              | 247        |
| 8.2               | Installing and using IBM Network Advisor . . . . .            | 248        |
| 8.2.1             | Installing IBM Network Advisor . . . . .                      | 249        |
| 8.2.2             | Using the IBM Network Advisor . . . . .                       | 259        |
| 8.3               | Setting up a FICON Director . . . . .                         | 262        |
| 8.3.1             | Changing the IP addresses . . . . .                           | 262        |
| 8.3.2             | Using the FICON wizard . . . . .                              | 267        |
| 8.3.3             | Enabling features (optional) . . . . .                        | 268        |
| 8.3.4             | Setting up a logical switch (optional) . . . . .              | 270        |
| 8.3.5             | Configuring the Domain ID and Insistent Domain ID . . . . .   | 275        |
| 8.3.6             | Setting up PBR, IOD, and DLS . . . . .                        | 277        |
| 8.3.7             | Enabling the Control Unit Port . . . . .                      | 279        |
| 8.3.8             | Changing the port type and speed . . . . .                    | 282        |
| 8.3.9             | Changing buffer credits . . . . .                             | 286        |
| 8.3.10            | Setting up the Allow/Prohibit Matrix (optional) . . . . .     | 288        |
| 8.3.11            | Setting up zoning (optional) . . . . .                        | 290        |
| 8.3.12            | Configuring Port Fencing (optional) . . . . .                 | 295        |
| 8.4               | Setting up a cascaded topology . . . . .                      | 298        |
| 8.4.1             | Setting up Inter-Switch Links . . . . .                       | 299        |
| 8.4.2             | Setting up a high integrity fabric . . . . .                  | 303        |
| 8.5               | FICON Directors in an extended distance environment . . . . . | 311        |
| 8.6               | FICON Directors in an intermixed environment . . . . .        | 313        |
| 8.7               | Channel swap . . . . .  | 313        |
| 8.8               | Backing up Director configuration data . . . . .              | 314        |
| 8.9               | Backing up IBM Network Advisor configuration data . . . . .   | 315        |

**Part 4. Managing the FICON environment . . . . . 317**

|                   |  |            |
|-------------------|--|------------|
| <b>Chapter 9.</b> | <b>Monitoring the FICON environment . . . . .</b>        | <b>319</b> |
| 9.1               | System Activity Display . . . . .                        | 320        |
| 9.2               | Resource Measurement Facility monitoring . . . . .       | 321        |
| 9.3               | Introduction to performance monitoring . . . . .         | 321        |
| 9.4               | Introduction to Resource Measurement Facility . . . . .  | 322        |
| 9.4.1             | Data gathering . . . . .                                 | 322        |
| 9.4.2             | RMF reporting . . . . .                                  | 322        |
| 9.5               | RMF example reports . . . . .                            | 323        |
| 9.5.1             | DASD Activity report . . . . .                           | 323        |
| 9.5.2             | I/O Queueing report . . . . .                            | 325        |
| 9.5.3             | Channel Path Activity report . . . . .                   | 326        |
| 9.5.4             | FICON Director Activity report . . . . .                 | 328        |
| 9.5.5             | Enterprise Disk Systems Link Statistics report . . . . . | 328        |
| 9.5.6             | General performance guidelines . . . . .                 | 330        |
| 9.5.7             | Tape . . . . .   | 331        |



|   |            |
|---|------------|
| 9.6 IBM Network Advisor performance monitor . . . . .                               | 332        |
| 9.6.1 Real-Time Performance Reporting . . . . .                                     | 332        |
| 9.6.2 Historical performance reporting . . . . .                                    | 335        |
| 9.6.3 End-to-End Monitoring . . . . .   | 342        |
| 9.6.4 Top Talker Monitoring . . . . .   | 343        |
| 9.7 Web Tools Performance Monitor . . . . .   | 346        |
| <b>Chapter 10. Debugging FICON problems . . . . .</b>                               | <b>349</b> |
| 10.1 Preparing for problem determination activities . . . . .                       | 350        |
| 10.1.1 Using the D M - Display Matrix command. . . . .                              | 350        |
| 10.1.2 Creating a CONFIG member . . . . .   | 350        |
| 10.1.3 IBM Health Checker: Command Response (CMR) Time Function. . . . .            | 351        |
| 10.1.4 Controlling planned and unplanned changes . . . . .                          | 351        |
| 10.1.5 Tracking some I/O messages . . . . .   | 352        |
| 10.2 Problem determination approach for FICON. . . . .                              | 353        |
| 10.3 Diagnostic procedure guideline. . . . .  | 355        |
| 10.4 Environmental record, editing, and printing program . . . . .                  | 357        |
| 10.5 FICON link incident reporting . . . . .  | 357        |
| 10.6 FICON Purge Path Extended . . . . .  | 357        |
| 10.7 Helpful z/OS commands . . . . .  | 360        |
| 10.8 Node descriptor. . . . .   | 362        |
| 10.8.1 View Node Descriptors from HMC/SE . . . . .                                  | 362        |
| 10.8.2 View Node Descriptors from IBM Network Advisor. . . . .                      | 362        |
| 10.9 IBM Network Advisor PD Tools. . . . .  | 364        |
| 10.9.1 Logs . . . . .   | 364        |
| 10.9.2 SupportSave . . . . .  | 365        |
| 10.9.3 Port Error Statistics. . . . .   | 368        |
| 10.9.4 Fibre Channel (FC) Troubleshooting. . . . .                                  | 370        |
| 10.9.5 Helpful Command Line Interface commands . . . . .                            | 375        |
| 10.10 HMC PD Tools . . . . .  | 377        |
| 10.11 Common z/OS FICON error message: Interpretation . . . . .                     | 377        |
| <b>Part 5. Appendixes . . . . .</b>   | <b>395</b> |
| <b>Appendix A. Example: Planning workflow . . . . .</b>                             | <b>397</b> |
| <b>Appendix B. Configuration worksheets . . . . .</b>                               | <b>405</b> |
| <b>Appendix C. Configuration and definition tools. . . . .</b>                      | <b>409</b> |
| Hardware Configuration Definition . . . . .   | 410        |
| Hardware Configuration Manager. . . . .   | 410        |
| CHPID Mapping Tool . . . . .  | 410        |
| I/O Configuration Program . . . . .   | 411        |
| Stand-Alone I/O Configuration Program . . . . .                                     | 411        |
| IOCP statements and keywords used for FICON . . . . .                               | 412        |
| <b>Appendix D. Configuring the DS8000 for FICON . . . . .</b>                       | <b>415</b> |
| DS8000 licenses for FICON. . . . .  | 416        |
| Configuring DS8000 I/O ports . . . . .  | 417        |
| I/O port layout . . . . .   | 417        |
| Configuring I/O ports using the DS8000 GUI . . . . .                                | 422        |
| <b>Appendix E. Using HMC and SE for problem determination information . . . . .</b> | <b>425</b> |
| E.1 HMC and SE information . . . . .  | 426        |
| E.1.1 HMC and SE versions and user IDs. . . . .                                     | 426        |

|  |            |
|--|------------|
| E.1.2 HMC and SE default user IDs and passwords . . . . .                | 427        |
| E.1.3 Logging on to the HMC and SE . . . . .                             | 427        |
| E.2 Displaying CPC resource information . . . . .                        | 429        |
| E.2.1 Displaying the System z CPC details panel . . . . .                | 429        |
| E.2.2 Finding a physical resource on a System z server . . . . .         | 431        |
| E.3 Displaying individual FICON channel information . . . . .            | 435        |
| E.3.1 FICON icon . . . . .   | 435        |
| E.3.2 Displaying detailed FICON channel information . . . . .            | 436        |
| E.3.3 Displaying the channel problem determination task panels . . . . . | 438        |
| E.4 Displaying FICON IFCC and Other errors . . . . .                     | 453        |
| E.4.1 Displaying the IFCC and Other panel . . . . .                      | 454        |
| E.4.2 FICON IFCC detected error . . . . .                                | 454        |
| E.4.3 FICON ELS (Extended Link Services) error detected error . . . . .  | 455        |
| E.4.4 FICON timeout error . . . . .                                      | 457        |
| E.4.5 FICON External error . . . . .                                     | 459        |
| E.4.6 Channel Hardware Failure . . . . .                                 | 461        |
| E.4.7 Purge Path Extended Log . . . . .                                  | 461        |
| E.5 Displaying the Fibre Channel Analyzer panel . . . . .                | 466        |
| E.6 Displaying the System I/O Configuration Analyzer panel . . . . .     | 469        |
| <b>Appendix F. Useful z/OS commands . . . . .</b>                        | <b>471</b> |
| Using z/OS commands for problem determination . . . . .                  | 472        |
| Displaying system status using D M=CPU . . . . .                         | 472        |
| Displaying additional z/OS information using D IPLINFO . . . . .         | 474        |
| Displaying the I/O configuration using D IOS,CONFIG . . . . .            | 475        |
| Displaying HSA usage using D IOS,CONFIG(HSA) . . . . .                   | 475        |
| Display units command D U . . . . .                                      | 476        |
| D U,,,dddd,1 . . . . .   | 477        |
| D M=DEV(dddd) . . . . .  | 480        |
| DEVSERV command - DS P,dddd,n . . . . .                                  | 484        |
| <b>Appendix G. Adding FICON CTC connections . . . . .</b>                | <b>489</b> |
| <b>Related publications . . . . .</b>                                    | <b>493</b> |
| IBM Redbooks . . . . .   | 493        |
| Other publications . . . . .   | 494        |
| Online resources . . . . .   | 495        |
| How to get Redbooks . . . . .  | 495        |
| Help from IBM . . . . .  | 495        |
| <b>Index . . . . .</b>   | <b>497</b> |

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

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

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](http://www.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:

|  |   |                 |
|--|---|-----------------|
| DB2®                                       | MVS™  | System z9®      |
| DS8000®                                    | Parallel Sysplex®   | System z®       |
| Easy Tier®                                 | PR/SM™  | Tivoli®         |
| ECKD™                                      | Redbooks®   | z/Architecture® |
| ESCON®                                     | Redbooks (logo)  ® | z/OS®           |
| FICON®                                     | Resource Link®  | z/VM®           |
| GDPS®                                      | Resource Measurement Facility™  | z/VSE®          |
| Geographically Dispersed Parallel Sysplex™ | RMF™  | z10™            |
| Global Technology Services®                | S/390®  | z9®             |
| HyperSwap®                                 | System Storage DS®  | zEnterprise™    |
| IBM®                                       | System Storage®   | zSeries®        |
|  | System z10®   |                 |

The following terms are trademarks of other companies:

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.

Snapshot, NOW, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Disk Magic, and the IntelliMagic logo are trademarks of IntelliMagic BV in the United States, other countries, or both.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

# Preface

This IBM® Redbooks® publication covers the planning, implementation, and management of IBM System z® FICON® environments. It discusses the FICON and Fibre Channel architectures, terminology, and supported topologies.

The book focuses on the hardware installation and the software definitions that are needed to provide connectivity to FICON environments. You will find configuration examples required to support FICON control units, FICON Channel-to-Channel (FCTC), and FICON Directors. It also discusses utilities and commands that are useful for monitoring and managing the FICON environment.

The target audience for this document includes IT Architects, data center planners, SAN administrators, and system programmers who plan for and configure FICON environments. You are expected to have a basic understanding of IBM zEnterprise™ and System z10® hardware, HCD, HCM or IOCP, and a broad understanding of the Fibre Channel and FICON architectures.

## The team who wrote this book

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

**Karan Singh** is a Project Leader at the IBM ITSO GCS organization in Poughkeepsie, New York.

**Marimar Grau** is a Senior IBM z/OS® IT Specialist in IBM Spain. She has 17 years of experience in IBM mainframe environment. She has worked at IBM for the last 12 years. She has experience in the z/OS Operating system, related software, and IBM Hardware solutions. She is currently working as a trusted advisor at one of the biggest financial institutions of Spain. She provides Technical support on IBM HW Solutions: IBM system z196 and z10, Brocade FICON Director, IBM Tape solution VE TS7740, and 3584 Tape Library.

**Peter Hoyle** works for IBM performing System z configuration in Australia. He has more than 30 years experience in the IT industry, including 20 years in IBM mainframe server configuration management and design. He has worked at IBM since 1999. His areas of expertise include HCM and HCD and configuration implementation solution design and documentation. He also co-authored IBM System z9® Configuration Setup, IBM System z10 Enterprise Class Configuration Setup, I/O Configuration using z/OS HCD and HCM, and IBM zEnterprise System Configuration Setup.

**Jonghak Kim** is a Senior Product Services-System Service Representative at the Global Technical Services in Seoul, Korea. He joined IBM in 1986 to support System z and Mainframe I/O products. He joined Technical Support Group (TSG) in 1997 as a System z second-level supporter and has led the local second-level support team for System z/p/i and Storage platforms in Korea since 2002. His main areas of expertise are System z Hardware, FICON Directors and Fiber Connectivity, and IBM Parallel Sysplex®.


**Jukka Myyryläinen** is a Senior IT Specialist in IBM Finland, providing storage implementation services, data migration services, and technical support to customers. He has 25 years of experience with IBM in the storage field. His areas of expertise include IBM

high-end disk and tape systems and SAN. He co-authored several IBM Redbooks publications on IBM enterprise storage systems. He has a Masters degree in Mathematics from the University of Helsinki.

**Mark Pope** works for IBM as a Hardware Support Specialist within the MTS High-Availability Domain in the UK. He joined IBM in 1976 and was a Large Systems CE until 1991 when he joined the Support Centre. His responsibilities are providing hardware support to CEs and customers on System z and High-Availability Storages Products within the UK.

**Jean Sylvester** joined IBM in 2002 in IOS/HCD/HCM Level 2 Support. She currently works on the IBM zSeries® microcode patch team in Poughkeepsie, NY, USA. She has a Masters degree in Computer Science from Marist College, in Poughkeepsie. Previously she co-authored another IBM Redbooks publication, *I/O Configuration using z/OS HCD and HCM*.

**Francisco Vilaca** joined IBM Brazil in 2009 as a System Specialist at IBM Global Technology Services®. He has 20 years of professional background in IT and telecommunication. He has a Bachelor of Science degree in Electrical Engineering, a Major in Telecommunication Engineering, and is highly experienced with SAN and Storage technologies. Francisco is a Brocade Certified Fabric Professional (BCFP), Brocade Certified Fabric Designer (BCFD), Brocade Certified SAN Manager (BCSM), Brocade Certified Fabric Administrator (BCFA), and Brocade Certified Architect for Ficon (BCAF).



*Figure 0-1 The authoring team: Jonghak Kim, Marimar Grau, Mark Pope, Jean Sylvester, Karan Singh, Francisco Vilaca, Jukka Myyryläinen, Peter A. Hoyle.*

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

Bob Haimowitz, Richard Conway, Bill White  
International Technical Support Organization, Poughkeepsie Center

Connie Beuselinck  
IBM System z Product Planning, Poughkeepsie

Charlie Hubert, Brian Jacobs, Sam Mercier  
IBM Vendor Solutions Connectivity (VSC) Lab, Poughkeepsie

Lou Ricci  
IBM Systems Software Development, Poughkeepsie

Jack Consoli  
Systems Engineer, Brocade Communications Systems, Inc.

Thanks to the authors of the previous editions of this book.

- ▶ Authors of the first edition, *FICON Implementation Guide*, published in February 2005, were:
  - Hans-Peter Eckam, IBM Germany
  - Iain Neville, IBM UK
- ▶ Authors of the second edition, *FICON Implementation Guide*, published in January 2006, were:
  - Hans-Peter Eckam, IBM Germany
  - Wolfgang Fries, IBM Germany
  - Iain Neville, IBM UK
- ▶ Authors of the third edition, *FICON Implementation Guide*, published in September 2009, were:
  - Bill White, IBM USA
  - Wolfgang Fries, IBM Germany
  - Brian Hatfield, IBM USA
  - Michal Holenia, IBM Slovakia
  - Dennis Ng, IBM USA
  - Ewerson Palacio, IBM Brazil
  - René Petry, IBM Germany

## 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](http://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](http://ibm.com/redbooks)

- ▶ Send your comments in an e-mail to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

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





# Part 1

## Understanding FICON

This part introduces FICON and explains how it is exploited by the System z channel, the FICON Director, and the control unit. It also provides information about the FICON and Fibre Channel architectures and their uses in System z environments.





# Introduction to FICON

The term Fibre Connection (FICON) represents the architecture as defined by the InterNational Committee for Information Technology Standards (INCITS) and published as ANSI standards. FICON also represents the names of the various System z server I/O features. In this chapter, we discuss the basic Fibre Channel (FC), FICON, High Performance FICON for System z (zHPF) terminology, and System z FICON support, benefits, operating modes, and topologies.

Throughout this chapter we use the term FICON to refer to FICON Express8S, FICON Express8, FICON Express4, FICON Express2, and FICON Express, except when the function being described is applicable to a specific FICON feature type.

**Note:** FICON Express2, FICON Express4, FICON Express8, and FICON Express8S features do not support FCV mode. FCV mode is available with FICON Express LX feature 2319 (carry forward only on System z10).

The System z10 is the last server family to support FICON Express and FICON Express2.

The zEnterprise 196 (z196) and zEnterprise 114 (z114) are the last systems to support FICON Express 4 features. Review the usage of your installed FICON Express4 channels and, where possible, migrate to FICON Express8S channels.

# 1.1 Basic Fibre Channel terminology

This section describes the Fibre Channel (FC) architecture and explains general terms that are used both in the FC and FICON environments.

One defining characteristic of Fibre Channel is that it allows many existing, well-known, and long-implemented channel and networking protocols to run over the same physical interface and media. Consequently, there is a rich set of standards and terminology associated with its architecture. FICON uses the Single-Byte Command Code Sets -3 (SB-3) and Single-Byte Command Code Sets -4 (SB-4) implementation within the Fibre Channel standard.

The FC architecture, as illustrated in Figure 1-1, is an integrated set of rules consisting of five layers (FC-0 through FC-4) for serial data transfer between computers, devices, and peripherals developed by the InterNational Committee for Information Technology Standards (INCITS) and the American National Standard Institute (ANSI).

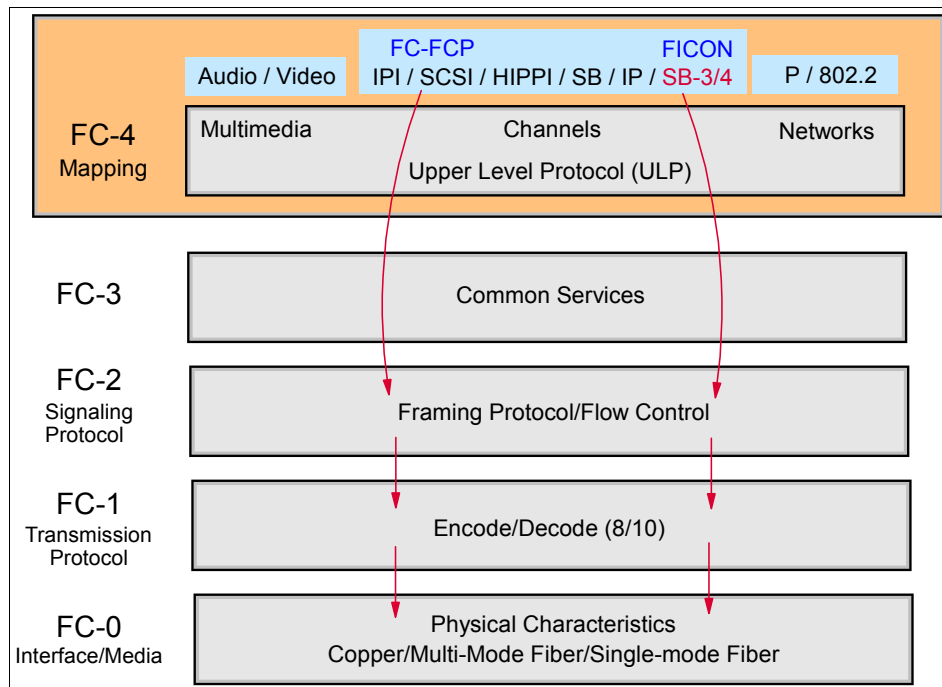


Figure 1-1 Fibre Channel architectural levels

The FICON channel architecture consists of the following Fibre Channel (FC) protocols:

- ▶ **FC-0 level: Interface and Media**

The Fibre Channel physical interface (FC-0), specified in FC-PI, consists of the transmission media, transmitters, receivers, and their interfaces. The physical interface specifies a variety of media and associated drivers and receivers that can operate at various speeds.

- ▶ **FC-1 level: Transmission Protocol**

This is a link control protocol that performs a conversion from the 8-bit EBCDIC code into a 10-bit transmission code. A unique bit-pattern is assigned to each known hexadecimal character. Encoding is done by the N\_Port when sending the character stream over the fiber, and the decoding back to 8-bit code is performed by the receiving N\_Port.

► FC-2 level: Signaling Protocol

Fibre Channel physical framing and signaling interface (FC-PH) describes the point-to-point physical interface, transmission protocol, and signaling protocol of high-performance serial links for support of higher-level protocols associated with HIPPI, IPI, SCSI, FC-SB-2/3/4 (FICON), and others.

► FC-3 level: Common Services

This layer is intended to provide the common services required for advanced features.

► FC-4 level: Protocol Mapping

The Upper Level Protocol (ULP) is part of FC-4 and describes IPI/FC-FCP (SCSI)/HIPPI/SB/IP and FICON:

- FICON was introduced with Single Byte-2 Command Code Sets mapping protocol, and then later revised with FC-SB-3. In 2008, another revision was drafted, identified as FC-SB-4. This revision is the current standard and is used to support additional FICON functions.

FC-SB-2, FC-SB-3, and FC-SB-4 architecture information and other FC documentation can be obtained from the following web site:

<http://www.t11.org>

Fibre Channel provides the capability to build a configuration, as shown in Figure 1-2, that can operate in a point-to-point, arbitrated loop, or switched fabric configuration.

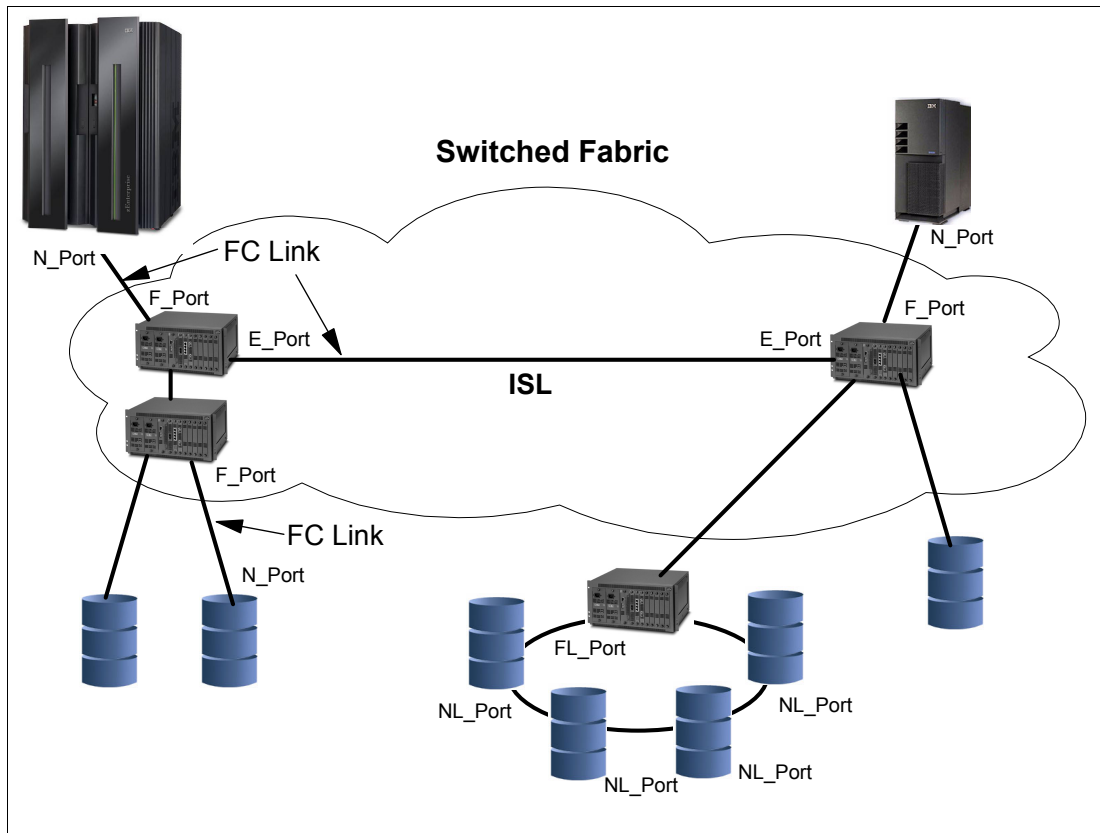


Figure 1-2 FC configuration with port types

**Note:** System z FICON channels do *not* support the arbitrated loop topology.

## 1.1.1 Node

A node is an endpoint for Fibre Channel (FC) traffic that contains information. It can be a computer (host), a device controller, or a peripheral device (such as a disk or tape drive). A node has a unique 64-bit identifier known as the Node\_Name, which is typically used for management purposes.

## 1.1.2 Port

Each node must have at least one port (hardware interface) to connect the node to the FC topology. This node port is referred to as an N\_Port. An N\_Port is not loop capable. They create, detect, and handle the flow of message units to and from the connected system.

Each N\_Port has a Port\_Name, which is a unique 64-bit identifier that is assigned at the time it is manufactured. The N\_Port associates an access point to a node's resources.

The storage area network (SAN) administrator assigns storage to a World Wide Port\_Name (WWPN), thus providing security to limit access to storage on the SAN.

FC standards allow N\_Port virtualization (NPIV) so that an N\_Port can be shared, requiring a unique WWPN to be generated for each NPIV.

Other port types include:

- E\_Port**      *An expansion port* interconnects switches and builds a switched fabric. E\_Ports can only attach to other E\_Ports.
- F\_Port**      *A fabric port* connects an N\_Port to a switch that is not loop-capable or for N\_Ports to access FC switched fabric functionality.
- FL\_Port**     *A fabric loop port* connects NL\_Ports to a switch in a loop configuration and is capable of Fibre Channel Arbitrated Loop (FC-AL) operations.
- G\_Port**      *A generic port* is a port that can operate as either an E\_Port or an F\_Port and this is determined at switch port initialization.
- L\_Port**      *A loop port* is a loop-capable fabric port in a FC-AL topology.
- NL\_Port**     *A node loop port* is an N\_Port operating on a loop.
- U\_Port**      *A universal port* is a port that has not yet assumed a specific function in the fabric. It is a generic switch port that can operate as an E\_Port, an F\_Port, or an FL\_Port.

The port type is determined by the node's role in the topology, as shown in Figure 1-2 on page 5.

## 1.1.3 Switched fabric

One or more switches are interconnected to create a fabric, to which the N\_Ports are connected. Switches can be interconnected with an Inter-Switch Link (ISL) to which the E\_Ports are connected. A switched fabric takes advantage of aggregated bandwidth through switched connections between N\_Ports and can offer better throughput than broadcast networks, such as Ethernet.

## 1.1.4 FC link

The port connects to the topology through an FC link. The FC link is a fiber optic cable that has two strands. It provides two dedicated, unidirectional, serial-bit transmission lines. One strand transmits a signal and the other strand receives a signal, as shown in Figure 1-3. An FC link interconnects nodes and switches. The Fibre Channel Standard (FCS) protocol specifies that for normal I/O operations frames flow serially in both directions, thereby allowing several concurrent read and write I/O operations on the same link.

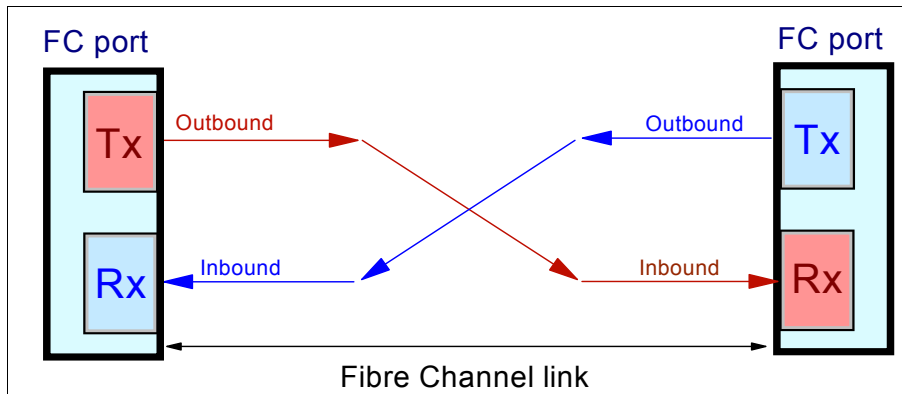


Figure 1-3 Fibre Channel link

An FC link (port-to-port connection) can be:

- ▶ Node-to-node link (N\_Port-to-N\_Port)
- ▶ Node-to-switch link (N\_Port-to-F\_Port)
- ▶ Switch-to-switch link (E\_Port-to-E\_Port)

## 1.1.5 World Wide Name

Nodes and ports have unique 64-bit addresses that identify them in an FC topology. These addresses are assigned by the manufacturer with a vendor-specific portion defined by the IEEE standards committee. These addresses (in the FC standard) are called Node\_Names and Port\_Names, and when they are worldwide unique, they are referred to as:

- ▶ World Wide Node\_Name (WWNN)
- ▶ World Wide Port\_Name (WWPN)

A WWN (any WWNN or WWPN) is usually written in sets of two hex digits, separated by colons (for example, 10:00:00:50:76:00:D5:A0). Figure 1-4 on page 8 illustrates where the WWNs are used.

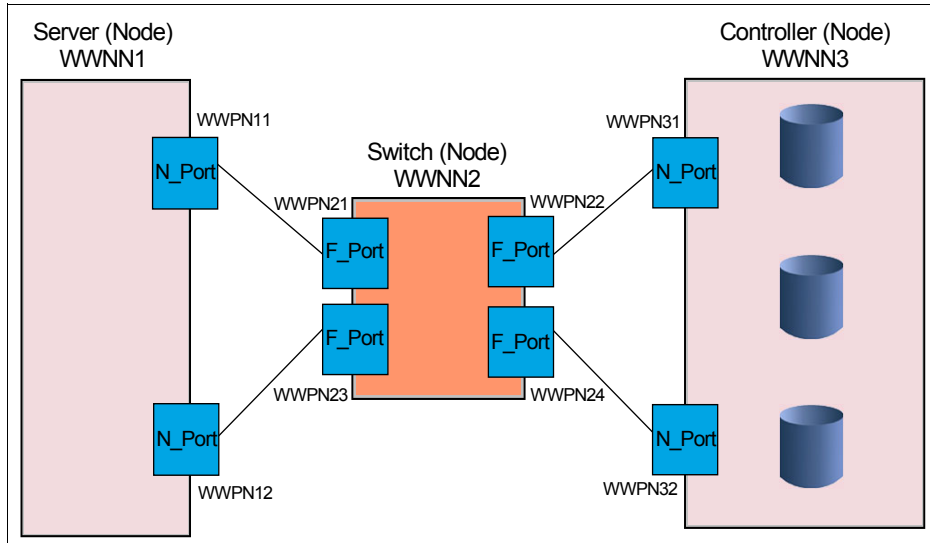


Figure 1-4 World Wide Names

On a System z server, the WWNN is constant for all FICON channels (ports); however, the WWPN is unique to each FICON channel on the server.

### 1.1.6 Open exchanges

An *open exchange* represents an I/O operation in progress over the FICON (or FC) channel. Many I/O operations can be in progress over FICON channels at any given time. As an example, a disk I/O operation might disconnect temporarily from the channel while waiting for a disk rotation or while performing a seek operation. Other I/O operations can be managed during this disconnect time, as follows:

- ▶ Command mode open exchanges

The number of open exchanges is limited by the FICON Express feature. FICON Express8S, FICON Express8, FICON Express4, and FICON Express2 allow up to 64 open exchanges. One open exchange (in reality, it is an exchange pair) in command mode is the same as one I/O operation in progress.

- ▶ Transport mode open exchanges

In transport mode, one exchange is sent from the channel to the CU. The same exchange ID is sent back from the control unit to the channel to complete the I/O operation. 750 exchanges is the maximum number of simultaneous exchanges that the channel can have open with a physical CU. The CU sets the maximum number of exchanges in the status area of the transport mode response Information Unit (IU).

**Note:** The default number of exchanges in transport mode is 64 and can be increased and decreased.

FICON channels can also multiplex data transfers for several devices simultaneously, allowing workloads with low to moderate control unit cache hit ratios to achieve higher levels of activity rates per channel.

If the open exchange limit is reached, the channel refuses additional I/O operations, which can result in queuing and retries by the operating system.



## 1.2 System z FICON

Because FICON is an IBM I/O technology built on the Fibre Channel architecture, many of the FC terms and definitions are used in the various FICON configurations.

FICON channel architecture is compatible with:

- ▶ Fibre Channel Physical and Signaling standard (FC-FS)
- ▶ Fibre Channel Switch Fabric and Switch Control Requirements (FC-SW)
- ▶ Fibre Channel Single-Byte-3 (FC-SB-3) and Fibre Channel Single-Byte-4 (FC-SB-4)

Even though FICON uses the FC architecture to transport commands and data over the fiber link, FICON is implemented differently from other FC implementations. For example, in a switched topology, a server using FC might use a discovery or definition process to obtain a port address. System z FICON only uses the definition process:

- ▶ Discovery

Uses the WWN and discovers the associated N\_Port address by querying the Fabric Name Server. This approach is used by all Fibre Channel protocol Host Bus Adapters (HBAs) users connected to a switched topology.

- ▶ Definition

Uses the Hardware Configuration Definition (HCD) tool, Hardware Configuration Manager (HCM) tool, or Input Output Configuration Program (IOCP) to define the associated N\_Ports to which communication is allowed. The definitions in HCD/HCM/IOCP must match the actual port connections of the attached control units.

FICON is widely used in the System z environment and provides additional strengths and capabilities compared to the IBM ESCON® technology. Many additional capabilities have been included in support of FICON since it was originally introduced. Some control units and control unit functions might require FICON use exclusively, for example, Hyper Parallel Access Volume requires the use of FICON and will not work with ESCON.

### ***Hyper Parallel Access Volume (HyperPAV)***

Hyper Parallel Access Volume (HyperPAV) is an on-demand automated allocation of alias devices for System z that enhances the original PAV feature supported on DASD subsystems. HyperPAVs can be used to improve I/O efficiency and performance, while simultaneously reducing the number of I/O resources used in System z environments. HyperPAVs require FICON connections. Additional support or licenses might be required on the control unit. Refer to *IBM System Storage DS8000: Architecture and Implementation*, SG24-8886 for more information about HyperPAV.

Other FICON enhancements might require the use of certain FICON features used on a specific System z server, such as the zEnterprise 196 (z196), zEnterprise 114 (z114), or System z10 servers.

I/O adapters or ports supporting FICON reside on control units, switches, and servers. The I/O adapter can support various FICON operating modes.

In this book we focus on the FICON native operating mode and the latest FICON capabilities available on the zEnterprise 196, zEnterprise 114, and System z10 servers.

## 1.2.1 High Performance FICON for System z (zHPF)

The zHPF is an enhancement of the FICON channel architecture and is compatible with:

- ▶ Fibre Channel Physical and Signaling standard (FC-FS)
- ▶ Fibre Channel Switch Fabric and Switch Control Requirements (FC-SW)
- ▶ Fibre Channel Single-Byte-4 (FC-SB-4) standards

You can reduce the FICON channel overhead by exploiting zHPF with the FICON channel, the z/OS operating system, and the control unit. You achieve this by protocol optimization by reducing the number of Information Units (IUs) processed, resulting in more efficient usage of the fiber link.

Both the existing FICON architecture and the zHPF architecture are supported by the FICON Express8S, FICON Express8, FICON Express4, and FICON Express2 features. In z/OS the existing FICON architecture is called *command mode*, and the zHPF architecture is called *transport mode*. Bit 13 in word 1 of the Operation Request Block (ORB) is used to determine whether the FICON channel is running in command or transport mode. The mode used for an I/O operation is based on the control unit (CU) supporting zHPF and settings in the z/OS operating system. The IECIOSxx parameter in SYS1.PARMLIB and SETIOS commands in z/OS can enable or disable zHPF. The D IOS,ZHPF system command now indicates whether zHPF is enabled, disabled, or not supported on the server.

```
IECIOSxx: ZHPF=YES|NO (permanent change to configuration)

SETIOS ZHPF=YES|NO (temporary change until next IPL or next SETIOS)

D IOS,ZHPF
  IOS630I hh.mm.ss FCX FACILITY | FIBER CHANNEL EXTENSIONS FACILITY
  IS ENABLED | DISABLED | NOT SUPPORTED
```

During link initialization, both the control unit and the channel indicate whether they support zHPF. In response to the Request Node ID (RNID) Extended Link Services (ELS), the Process Login (PRLI) support indicator is presented. If PRLI is supported, the channel sends a PRLI ELS. The PRLI response indicates that zHPF is supported by the CU. As with the existing FICON channel architecture, the application or access method provides the channel program (CCWs) and parameters in the ORB.

Table 1-1 indicates that a device is enabled for zHPF when all of the qualifications are true.

Table 1-1 Devices are enabled for zHPF capability

| Qualifications   |
|--|
| 1. Processor supports zHPF   |
| 2. All of the online paths support zHPF <ul style="list-style-type: none"> <li>▶ CONFIG CHP(xx), ONLINE/OFFLINE and VARY PATH commands can affect status</li> </ul>  |
| 3. Control unit supports zHPF <ul style="list-style-type: none"> <li>▶ PSF/RSSD/Read Feature Codes - zHPF indicator</li> </ul>   |
| 4. Device support code supports zHPF <ul style="list-style-type: none"> <li>▶ Supported only for DASD (through UIM, DFPCX flag)</li> </ul>   |
| 5. The customer enabled zHPF in z/OS <ul style="list-style-type: none"> <li>▶ Using the ZHPF=YES/NO option in IECIOSxx parmlib or SETIOS command; default is zHPF is disabled.</li> <li>▶ UCBFCX bit is set if the device is enabled for zHPF</li> </ul> |

zHPF (transport mode) manages CCW operation significantly different from the existing FICON architecture's CCW operation (command mode). In command mode, each single CCW is sent to the control unit for execution, while in transport mode an I/O operation will use a transport command word (TCW) to send all the commands and data at one time to the control unit in a single frame instead of the multiple CCWs separately. Certain complex CCW chains are not supported by zHPF.

Figure 1-5 shows an example of how zHPF's transport mode optimizes a read operation.

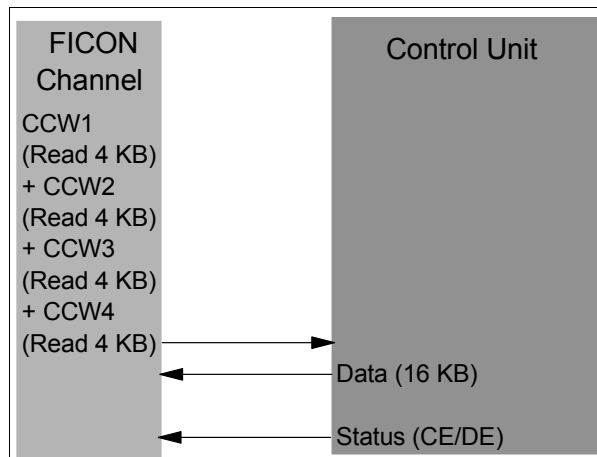


Figure 1-5 High-performance FICON read operation

All of the required CCWs and read operations of 4 KB of data are sent by the channel in one single frame to the CU. The requested data gets transferred over the link to the channel by the CU, followed by a CE/DE when the operation is successful, which generates less overhead than the existing FICON architecture.

Figure 1-6 shows the same reduction of frames and open exchanges for zHPF's write operation.

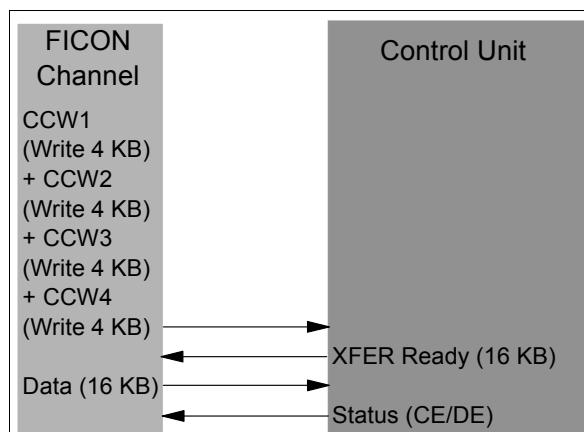


Figure 1-6 High-performance FICON write operation

The channel sends all the required CCWs and write operations of 4 KB of data in one frame to the CU. The CU responds with a XFER ready when it is ready to receive the data. The channel then sends the 16 KB of data (in this case) to the CU. If the CU successfully receives the data and finishes the write operation, the CE/DE status is sent by the CU to indicate the completion of the write operation.

zHPF supports multitrack operations allowing the channel to operate at rates that fully exploit the bandwidth of a FICON Express channel. The zHPF fully supports multiple tracks of data that can be transferred in a single operation.

**Note:** All FICON channel paths (CHPIDs) defined to the same Logical Control Unit (LCU) must support zHPF. Including any non-compliant zHPF features in the path group causes the entire path group to support command mode only.

## High performance FICON for System z enhancements

z196 and z114 with zHPF offer improved access to data for Online transaction processing (OLTP) applications. They also offer 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 FCP.

With IOS/EXCP APAR OA 33642 and DFSMS Device Support APAR OA34662, zHPF now handles QSAM, BPAM, and BSAM access methods.

With IOS APAR OA33089 and OA34149 and DFSMS Device Support APAR OA34662 and DFSMS Media Manager APAR OA34661, zHPF also offers improved IBM DB2® list prefetch for certain DB2 queries and some DB2 utility operations.

### ***zHPF now for QSAM, BPAM, and BSAM access methods***

Existing zHPF support for VSAM, Extended Format sequential, zSeries File System (zFS), and PDSE data sets is changed to provide support for QSAM, BPAM, and BSAM data set types when a new parameter is specified in the IGDSMSxx parmlib member. Significant I/O performance improvements are expected without the need for application changes.

The zHPF enhancement for QSAM, BPAM, and BSAM access methods is exclusive to z196 and z114 and applies to all supported FICON features (CHPID type FC).

### ***zHPF DB2 list PREFETCH and FICON Express8S environments***

The zHPF enhancement to DB2 list prefetch processing is expected to provide significant performance improvements for certain DB2 queries and some DB2 utility operations.

DB2 uses prefetching to read multiple records from a data base asynchronously, in anticipation of using them for a query. Each record is 4K in size. One type of prefetch is called sequential prefetch and is used to read 128K or 256K of contiguous records on disk. Another is called list prefetch and is used to read 128K of non-contiguous records on disk. Sequential prefetch is supported by the control unit today with the IBM DS8000® multi-track feature. List prefetch is now supported through a combination of the new imbedded locate record and zHPFs bidirectional features.

The zHPF architecture limits the size of the Transport Command Area (TCA) command block to 240 bytes meaning that the prefix command is limited to a maximum of 22 non-contiguous records, which is not enough to satisfy DB2 list prefetch requests (DB2 needs 32). The zHPF channel program only allows command information (the Transport Command Control Block's TCA) to be sent in the first phase of the channel program (the command phase) because of the FCP protocol used. The remaining phases transfer data (read or write data or both).

The Transport Command Control Block (TCCB) was extended by transporting additional command information in the output stream. This is now a bidirectional request because the channel program is reading data but transporting additional command information as write data. Therefore, data transfer occurs in both directions (it does not matter what is in the data).

**Note:** z/OS rarely has a need to do true bidirectional transfer. Most channel programs either read or write data, not both.

The extended command block is called a Transport Command Area Extension (TCAX) and is similar to the TCA. It consists of one or more Device Command Words (DCWs) optionally followed by control data. The control data from a DCW within the TCA cannot be continued into the TCAX because the TCAX must always start with a DCW. The size of the TCAX is limited to 4K by the control unit, but architecturally up to 64K can be supported.

Table 1-2 indicates that a device has zHPF bidirectional (BiDi) capability when all of the following qualifications are true:

*Table 1-2 Devices have zHPF bidirectional (BiDi) capability*

| Qualifications   |
|--|
| 1. All qualifications in Table 1-1 on page 10 are met.   |
| 2. Control unit supports zHPF BiDi.  |
| 3. All online paths (channels) to the device support zHPF BiDi: <ul style="list-style-type: none"><li>▶ New UCBX_zHPF_BiDi bit is set on when all the online channels for the device support bidirectional data transport.</li><li>▶ z/OS (IOS) only determines the device's zHPF bidirectional capability from a processor point of view.</li></ul> |

The zHPF enhancement for DB2 list prefetch is exclusive to z196 and z114 and only applies to the FICON Express8S features (CHPID type FC).

### **IPL from an alternate subchannel set**

IPL from an alternate subchannel set is exclusive to z196 and z114. Multiple-Subchannel-Set facility must be enabled by the Hardware Management Console (HMC) when an IPL is done with a device in an alternate subchannel set. It is supported by z/OS V1.13, in addition to V1.12 and V1.11 with PTFs and applies to both the FICON and zHPF protocols.

On z196 and z114 you can IPL from subchannel set 1 (SS1), and on z196 only, you can IPL from subchannel set 2 (SS2) in addition to subchannel set 0 (SS0). Devices that are used early in IPL processing can now be accessed using SS1 (z196 and z114) and SS2 (z196 only). This 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, IODF, and stand-alone dump volumes when needed.

**Note:** Subchannel set 2 (SS2) is only available on z196.

## **1.2.2 FICON-to-ESCON solutions**

The FICON Express feature was the last FICON channel feature to support CHPID type FCV. The FICON Express channel in FCV mode allows the accessing of ESCON interface control units with ESCON interfaces using a 9032-5 ESCON Director FICON Bridge adapter. FCV mode is available only on a System z10 that has gone through a server upgrade.

The z196 and z114 are the last systems to support FICON Express4 features. The System z10 is the last server family to support FICON Express and FICON Express2. The FICON Express features are supported on z10 servers only if carried forward on a server upgrade.

Therefore, another solution is needed for FICON channel-to-ESCON control unit connectivity, such as the PRIZM FICON to ESCON Converter from Optica Technologies Inc.

The PRIZM FICON to ESCON Converter is a channel-based appliance that converts native FICON protocol to native ESCON protocol, thereby allowing attachment of existing ESCON devices directly to FICON channels.

Unlike the 9032-5 ESCON Director FICON Bridge solution, PRIZM supports native FICON channel protocol (CHPID type FC), and supports attachment to existing FICON fabrics.

For zHPF channel programs to support the Optima PRIZM FICON to ESCON converter they must have z/OS APAR OA29017 installed or above.

PRIZM provides strategic flexibility to those migrating from ESCON to FICON on their System z servers. For additional details, refer to the following web site:

[http://www.opticatech.com/?page\\_id=19](http://www.opticatech.com/?page_id=19)

The PRIZM solution is also offered through our IBM Partner web site:

<http://www-304.ibm.com/partnerworld/gsd/solutiondetails.do?&solution=26320&lc=en>

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

### 1.2.3 FICON operating modes and topologies

System z FICON features can be defined in HCD, HCM, or IOCP to operate in one of up to three operating modes. The defined operating mode supported depends on the FICON feature installed in the System z server. The three operating modes that can be defined are:

- ▶ FICON conversion mode (FCV): This operating mode is only supported on the FICON Express LX feature.
- ▶ FICON native mode (FC): This operating mode is for FICON, zHPF, and FICON channel-to-channel (FCTC) traffic for z/OS, IBM z/VM®, IBM z/VSE®, z/TPF, and Linux on System z environments.
- ▶ Fibre Channel Protocol (FCP): This operating mode is to attach SCSI devices for z/VM, z/VSE, and Linux on System z environments.

As shown in Figure 1-7 on page 15, a FICON channel in native mode can access FICON control units in any of the following three topologies:

- ▶ Point-to-point (direct connection)
- ▶ Switched (through a FICON Director)
- ▶ Cascaded (through *two* FICON Directors)

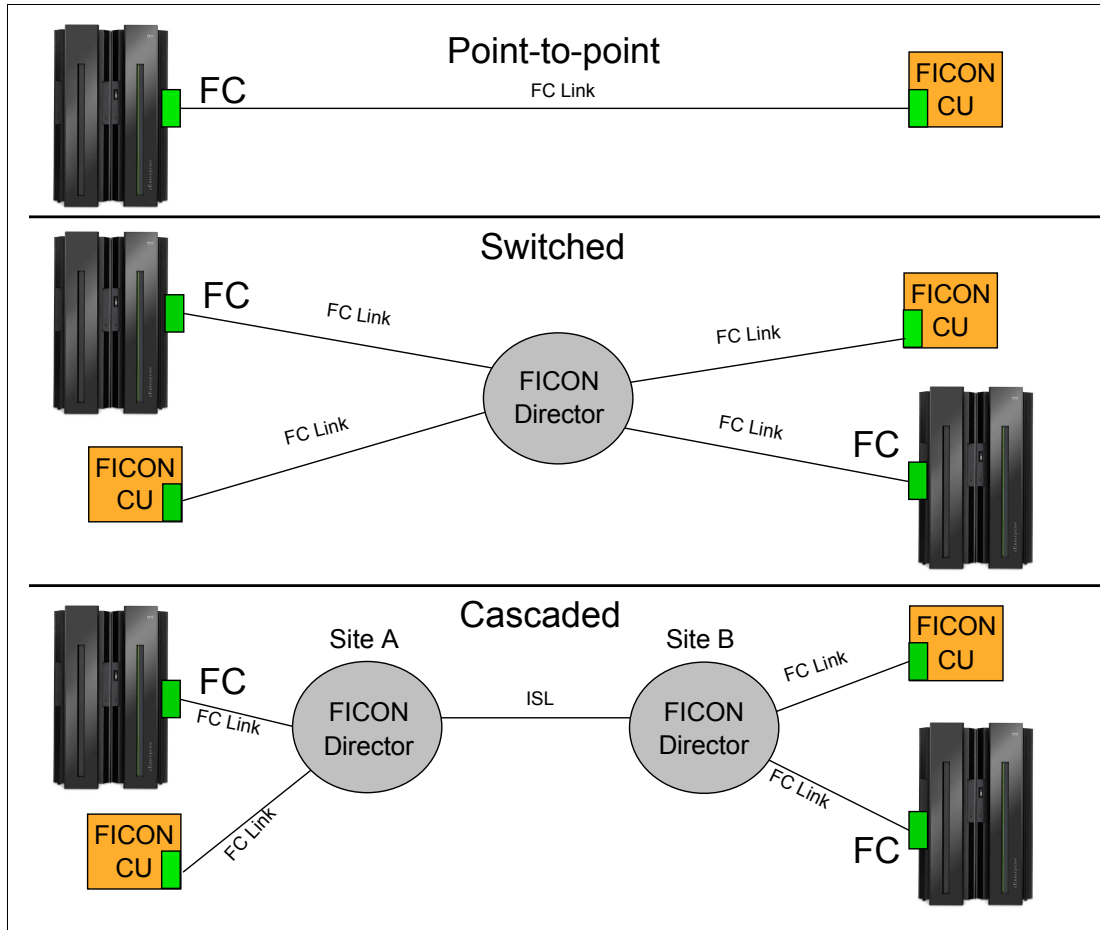


Figure 1-7 System z supported FICON native topologies

A FICON channel also supports channel-to-channel (CTC) communications. The FICON channel at each end of the FICON CTC connection, supporting the FCTC control units, can also communicate with other FICON control units, such as disk and tape devices.

## 1.2.4 Platform and name registration in FICON channel

All FICON features on the z196, z114, and z10 support platform and name server registration to the fabric when the FICON feature is defined as CHPID type FC.

If registered, the information about the channels connected to a fabric allows other nodes or storage area network (SAN) administrators to query the name server to determine what is connected to the fabric. The registered attributes for the z196, z114, and z10 are:

- ▶ Platform information:
  - World Wide Node Name (WWNN): This is the platform's node name and is the same for all channels belonging to the platform (see Figure 1-4 on page 8).
  - Platform type.
  - Host computer type.
  - Platform name: This includes vendor ID, product ID, and vendor-specific data from the node descriptor.

- ▶ Channel information:
  - World Wide Port Name (WWPN)
  - Port type (N\_Port\_ID)
  - FC-4 types supported
  - Classes of service supported by the channel

The platform and name server registration service are defined in the Fibre Channel - Generic Services 4 (FC-GS-4) standard (INCITS (ANSI) T11 group).

## 1.2.5 Terms used with FICON Directors

It is important to point out that some definition terms used in IOCP, HCD, and HCM were carried over from the ESCON environment. For clarity, we use the following terms throughout this book to describe the definitions needed in IOCP, HCD, and HCM for cascaded FICON Directors and switched configurations:

- ▶ An entry switch is the FICON Director that is directly connected to the processor's FICON channel and to the CU (destination) or another FICON Director.
- ▶ A cascaded switch is the FICON Director that connects to the CU (destination) and to the entry switch.
- ▶ The entry switch and cascaded switch are interconnected through an Inter-Switch Link (ISL).
- ▶ Switch ID and switch address (1-byte value) are terms that address a FICON Director.
- ▶ A port address (1-byte value) addresses the physical port on the FICON Director.
- ▶ The switch control unit port (CUP) function allows in-band management of a FICON Director through the FICON channel with the same level of control and security as for an ESCON switch.

Host communication includes control functions, such as blocking and unblocking ports, as well as monitoring and error reporting functions.

## 1.2.6 Terms used with the Input/Output architecture

The System z channel environment uses many terms and acronyms. We define some of those terms here. The terms listed here are used throughout this book as we discuss channel and I/O operations in the various chapters.

### ***Input/output channels***

Input/output (I/O) channels are components of the System z Channel Subsystem (CSS) and IBM z/Architecture®. They provide a pipeline through which data is exchanged between servers, or between a server and external devices. z/Architecture channel connections are referred to as *channel paths*.

### ***Control unit***

The most common attachment to a System z channel is a control unit (CU) accessed using an Enterprise Systems CONnection (ESCON) or a Fibre CONnection (FICON) channel. The CU controls I/O devices, such as disk and tape devices.

### ***I/O devices***

An input/output (I/O) device provides external storage, a means of communication between data processing systems, or a means of communication between a system and its



environment. In the simplest case, an I/O device is attached to one control unit and is accessible through one channel path.

### **Channel-to-Channel**

The Channel-to-Channel (CTC) function simulates an I/O device that can be used by one system control program to communicate with another system control program. It provides the data path and synchronization for data transfer between two channels. When the CTC option is used to connect two channels that are associated with different systems, a loosely coupled multiprocessing system is established. The CTC connection, as viewed by either of the channels it connects, has the appearance of an unshared input/output device.

### **Channel Subsystem**

The Channel Subsystem provides the functionality for System z servers to communicate with input/output (I/O) devices and the network. The CSS evolved with the increased scalability of IBM System z servers.

The CSS architecture provides functionality in the form of multiple Channel Subsystems (CSSs). Multiple CSSs can be configured within the same System z server. Each CSS can have up to 256 channels and, in turn, can be configured with up to 15 logical partitions.

Table 1-3 lists the maximum number of CSSs or CSS image IDs (numbered 0 to 3) and logical partitions supported by System z servers. System z servers are designed to offer considerable processing power, memory sizes, and I/O connectivity. For the System z with larger I/O capability, the CSS concept was scaled up accordingly to provide relief for the number of supported logical partitions, channels, and devices available to the server.

*Table 1-3 Maximum number of CSSs and logical partitions supported by System z servers*

| <b>Servers</b> | <b>Maximum number of CSSs</b> | <b>Maximum number of logical partitions</b> |
|----------------|-------------------------------|---|
| z196           | 4                             | 60  |
| z10 EC         | 4                             | 60  |
| z114           | 2                             | 30  |
| z10 BC         | 2                             | 30  |

### **Subchannels**

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

Multiple Subchannel Sets (MSS) are available on z196, z114, and z10 to increase addressability. Three subchannel sets are provided: subchannel set-0 (SS0) can have up to 63.75 K subchannels, and subchannel set-1 (SS1) and subchannel set-2 (SS2) can have up to 64 K subchannels. Subchannel set-2 (SS2) is available on z196 only.

SS1 and SS2 both can be used for disk alias devices of both primary and secondary devices. You must include the Subchannel Set number in HCD/HCM/IOCP definitions. The default is SS0.

### **Channel Path Identifier**

A Channel Path Identifier (CHPID) is a value assigned to each channel path of the system that uniquely identifies that path. A total of 256 CHPIDs are supported by each CSS.

### ***Multiple Image Facility***

Multiple Image Facility (MIF) enables resource sharing across logical partitions within a single CSS or across CSSs.

### ***Channel spanning***

Spanning is the ability to configure channels to multiple Channel Subsystems. This way the channels can be transparently shared by any or all of the configured logical partitions, regardless of the channel subsystem to which the logical partition is configured.

FICON channels can be spanned across multiple CSSs in z196, z114, and z10 and the CHPID number is the same on all logical CSSs that include that CHPID.

### ***Control Unit Port (CUP)***

The CUP function allows in-band management of a FICON Director through the FICON channel with the same level of control and security as for an ESCON switch. Host communication includes control functions like blocking/unblocking ports, performance monitoring, and error reporting functions.

### ***IBM Tivoli System Automation for z/OS (SA for z/OS)***

IBM Tivoli® System Automation for z/OS (SA for z/OS) includes support for FICON channels and FICON Directors. You can find additional information, updates, extensions, tools, and demonstrations for SA for z/OS at the following web site:

<http://www.ibm.com/software/tivoli/products/system-automation-390>

Before using SA for z/OS in your FICON environment, check the latest maintenance recommendations in the appropriate z/OS subset of the 2817DEVICE, 2818DEVICE, 2097DEVICE, and 2098DEVICE Preventative Service Planning (PSP) buckets prior to implementation.

## **1.3 z/OS discovery and auto-configuration (zDAC)**

z/OS discovery and auto-configuration (zDAC) is a function, exclusive to z196 and z114 (with FICON Express8S, FICON Express8, and FICON Express4), that automatically performs a number of I/O configuration definition tasks for new and changed FICON disk and tape control units that are connected to a FICON Director. It is designed to help simplify I/O configuration of zEnterprise CPCs running z/OS, and to help reduce complexity and setup time. We suggest using FICON Dynamic Channel Management (DCM in z/OS) to help manage performance.

The zDAC function is integrated into the existing Hardware Configuration Definition (HCD) and Hardware Configuration Manager (HCM) tools. A policy can be defined in HCD/HCM according to the availability and bandwidth requirements, including parallel access volume (PAV) definitions, control unit numbers, and device number ranges. When new controllers are added to an I/O configuration or changes are made to existing controllers, the system is designed to discover them and propose configuration changes based on that policy. The zDAC proposed configurations are created as work I/O definition files (IODF) that can be converted to production IODF and activated.

zDAC provides real-time discovery for the FICON fabric, subsystem, and I/O device resource changes from z/OS. By exploring the discovered control units for defined logical control units (LCU) and devices, zDAC compares the discovered controller information with the current system configuration to determine delta changes to the configuration for a proposed configuration. All newly added or changed LCUs and devices are added into the proposed

configuration with proposed control unit and device numbers and channel paths based on the defined policy.

zDAC uses a channel path chosen algorithm to minimize single point of failure. zDAC applies to all FICON features supported on z196 and z114 when configured as CHPID type FC.

To start zDAC you must have UPDATE authority to IBM MVS™.ACTIVATE OPERCMDS resource. New IOS530I, IOS531I, and IOS532I messages are issued through HCD/HCM panels similar to IOS500I messages. Explanations and actions are only documented in HCD/HCM help panels.





# System z FICON technical description

FICON uses the Fibre Channel Single Byte Command Sets-3 (FC-SB-3) standard as defined by the InterNational Committee for Information Technology Standards (INCITS) and published as ANSI standards. zHPF uses the architecture as described in the Fibre Channel Single Byte Command Sets-4 (FC-SB-4). FC-SB-4 is intended to be a complete replacement of the FC-SB-3 standard.

We provide a technical description of how System z servers use the FICON architecture to process I/O operations, as well as the FICON features and functions available on zEnterprise 196 (z196), zEnterprise 114 (z114), and System z10 servers.

## 2.1 Using the FICON architecture for I/O operations

Before a System z FICON channel can send any I/O over the link, the link must be initialized. There are specific steps used in this initialization process, and they are described in detail in the next section.

After the initialization process is complete, I/O operations can be started on the FICON channel. A FICON channel operates in what is known as *command mode* and in *transport mode* (also known as zHPF) simultaneously, exploiting the FC-SB-4 protocols. Note the following points:

- ▶ When both the FICON channel and the control unit indicate support for zHPF, the channel performs the I/O operation in transport mode.
- ▶ If the control unit does not support zHPF, the FICON channel operates in command mode.
- ▶ zHPF must be enabled in z/OS to perform transport mode operations.

Both modes (command and transport) are discussed in greater detail later in this chapter.

### 2.1.1 FICON initialization

The FICON channel initialization process establishes the necessary conditions for elements of a channel path to be able to sustain both link-level and device-level communication. Under normal conditions, the initialization process occurs infrequently, for example, it is usually required only after a power-on procedure or during a system initialization procedure, and is *not* part of the execution of an I/O operation. There is a set of initialization steps that is used with every connection. There is also a series of steps used for every control unit image stored in the HSA. For example, the initialization process for a FICON channel consists of the following procedures:

- ▶ Link initialization
- ▶ Channel login and security attribute determination
- ▶ Channel node-identifier acquisition
- ▶ Channel state-change registration
- ▶ Channel link-incident-record registration
- ▶ Process login, when supported by the channel and control unit
- ▶ Channel logical-path establishment

For a control unit, the initialization process consists of the following procedures:

- ▶ Link initialization
- ▶ Control unit login
- ▶ Control unit node-identifier acquisition
- ▶ Control unit state-change registration
- ▶ Process login, if supported
- ▶ Control unit logical-path establishment

The Fabric login (FLOGI) and Port login (PLOGI) process determine the presence of a switch fabric. If a switch fabric is present, the channel performs an F\_Port Login. If a switch fabric is *not* present, the channel performs an N\_Port Login.

#### Point-to-point

This is the initialization process for a point-to-point connection:

1. Channel F\_Port Login process (to determine if connected to a switch).
2. Channel N\_Port Login process.

3. Request Node Identifier (RNID) function to provide specific neighbor node information.<sup>1</sup>
4. Channel State-Change Registration (SCR) is sent to the control unit.
5. Channel Link-Incident-Record Registration (LIRR).
6. Process Login (PRLI) is used when channel and control unit support System z High Performance FICON (zHPF).
7. Establish Logical Path (ELP) to the control unit images that are stored in the Hardware Storage Area (HSA<sup>2</sup>) when each channel image is initialized.

## Switched

This is the initialization process for a switched connection:

1. Channel F\_Port Login process.
2. Query Security Attributes (QSA) determines if a cascaded switch configuration is supported and if two-byte destination addresses can be used.
3. RNID function provides specific neighbor node information.
4. Channel SCR is sent to the fabric controller.
5. LIRR channel to the management server.
6. Channel N\_Port Login.

The channel will also login to the defined control unit N\_Port. The channel performs this for each control unit N\_Port link address defined on this channel path:

- RNID to control unit.
- LIRR to control unit.
- PRLI used when channel and control unit support zHPF.
- ELP to the control unit images that are stored in the HSA.

## Cascaded

The cascaded FICON Directors login process is similar to the process of the switched process. The only difference is that the FICON channel and control unit log into (F\_Port Login) two different FICON Directors.

Figure 2-1 shows the key steps used in the FICON initialization process for a switched connection with a control unit that supports zHPF.

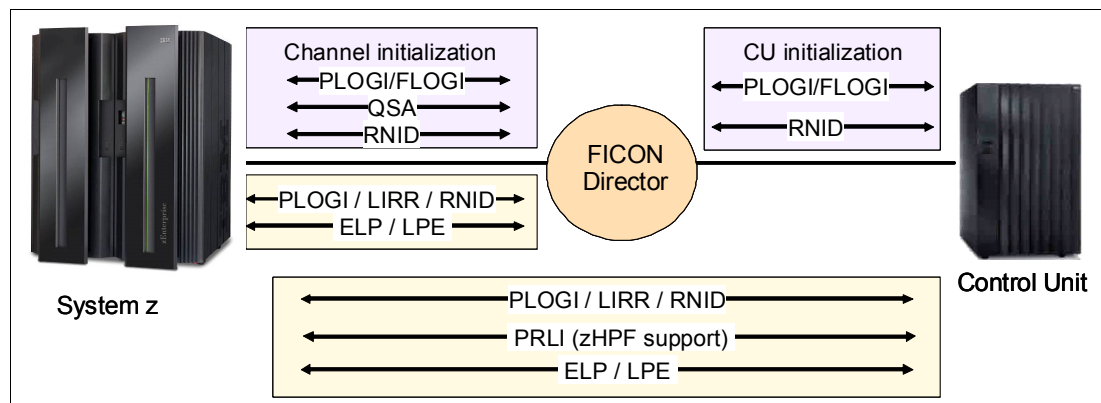


Figure 2-1 FICON initialization - channel to control unit initialization

<sup>1</sup> Request node identification data (RNID) function for native FICON CHPID type FC also allows isolation of cabling-detected errors.

<sup>2</sup> Also called the Hardware System Area, it is a protected area of storage where control blocks used by the channel subsystem are located.

Link initialization is described in FC-FS-2. When link initialization is complete, the N\_Port or F\_Port is in the active state. After link initialization is complete for an N\_Port or F\_Port, the port is considered to be operational as long as it remains in the active state.

The link speed is also determined during this process, through auto-negotiation. The FICON link will start at the highest speed and work lower until both sides choose the same speed. FICON features support the following speeds:

- ▶ FICON Express2 will auto-negotiate to 2 Gbps or 1 Gbps.
- ▶ FICON Express4 will auto-negotiate to 4 Gbps, 2 Gbps, or 1 Gbps.
- ▶ FICON Express8 will auto-negotiate to 8 Gbps, 4 Gbps, or 2 Gbps.
- ▶ FICON Express8S will auto-negotiate to 8 Gbps, 4 Gbps, or 2 Gbps.

### **FICON registration**

z196, z114, and z10 servers support platform and name server registration to the fabric. When registered, information about the channels connected to a fabric is available to other nodes or SAN managers. The attributes that are registered for the System z servers are:

- ▶ Platform information:
  - World Wide Node Name (WWNN): This is the node name of the platform, and it is the same for all channels belonging to the platform.
  - Platform type.
  - Host computer type.
  - Platform name: The name includes Vendor ID, product ID, and vendor-specific data from the node descriptor.
- ▶ Channel information:
  - World Wide Port Name (WWPN).
  - Port type (N\_Port\_ID).
  - FC-4 types supported.
  - Classes of service supported by the channel.

Platform registration is a service defined in the Fibre Channel - Generic Services - 4 (FC-GS-4) standard (INCITS (ANSI) T11 group).

## **2.1.2 FICON I/O request**

Before the actual I/O operation starts and FC frames are transmitted over the fiber cable, the channel program must be created and passed to the Channel Subsystem (CSS). The Input/Output Supervisor (IOS) is a z/OS component, between the CSS and the application program, that issues the start subchannel command. The application program uses an access method that calls the I/O driver to build the channel program. The channel program might use transport commands when operating in transport mode or channel commands when operating in command mode.

- ▶ Transport mode is used by the System z High Performance FICON (zHPF) architecture. A transport mode I/O operation uses a Transport Control Word (TCW) to send commands and data over the FC link and does not use CCWs. The channel and control unit both must support zHPF.
- ▶ Command mode is used by the existing FICON architecture where an I/O operation consists of CCWs and data.



When the SSCH is passed to the CSS, as shown in Figure 2-2, and the FICON channel is selected, the FC architecture and FICON protocols are used to build the frames required to send over the FC link.

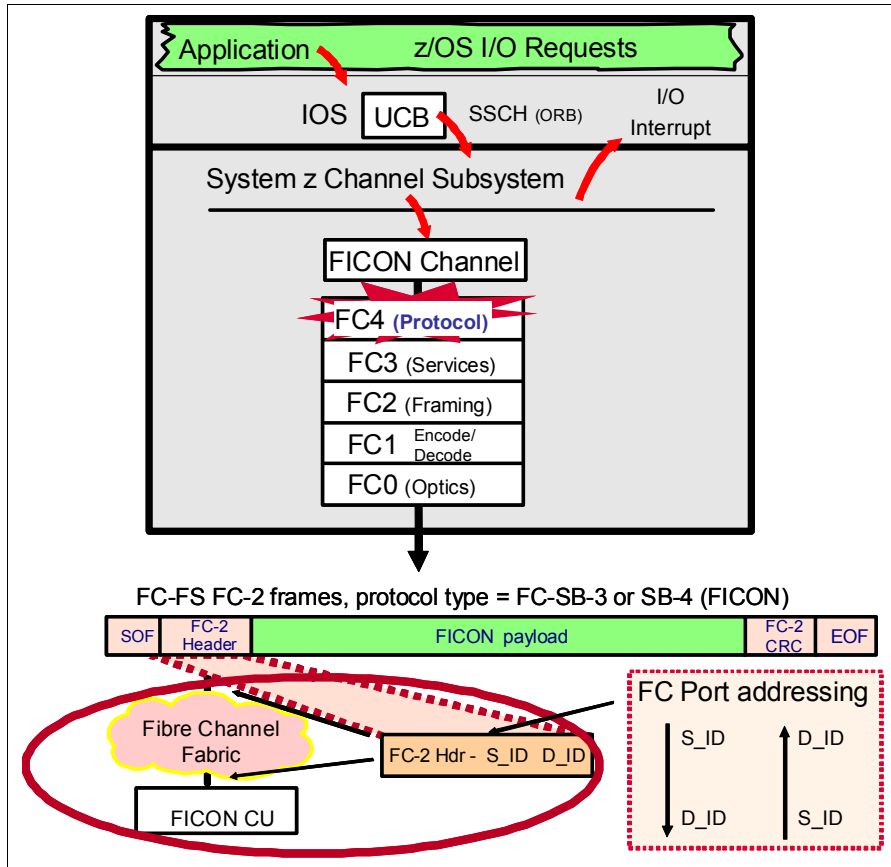


Figure 2-2 FICON channel I/O request

A FICON I/O request flow is as follows:

1. An application or system component invokes an I/O request. The application or access method provides CCWs or TCWs and additional parameters in the Operation Request Block (ORB).
2. The request is queued on the Unit Control Block (UCB)<sup>3</sup>. The Input Output Supervisor (IOS) services the request from the UCB on a priority basis.
3. IOS issues a Start Subchannel (SSCH) instruction with the Subsystem Identification word (SSID) representing the device and ORB as operands.
4. The ORB contains start-specific control information. It indicates whether the channel is operating in transport mode (zHPF support) or command mode, and also indicates the starting address of the Channel Programs (CPA) Channel Command.
5. The Channel Subsystem (CSS) selects the most appropriate channel and passes the I/O request to it.
6. The channel fetches from storage the Channel Command Words (CCWs) or Transport Command Words (TCWs) and associated data (for Write Operations).

<sup>3</sup> The UCB is a control block in memory that describes an I/O device to the operating system.

7. The channel assembles the required parameters and fields of the FC-2 and FC-SB-3 or FC-SB-4 for the I/O request and passes them to the Fibre Channel adapter (which is part of the FICON channel).

Device-level information is transferred between a channel and a control unit in SB-3 or SB-4 Information Units (IUs). Information units are transferred using both link-level and device-level functions and protocols, for example, when the channel receives an initiative to start an I/O operation, the device-level functions and protocols obtain the command and other parameters from the current CCW or TCW and insert them into the appropriate fields within a command IU. When the command IU is ready for transmission, link-level functions and protocols provide additional information (for example, address identifiers and exchange ID in the frame header) and then coordinate the actual transmission of the frame on the channel path.

8. The Fibre Channel adapter builds the complete FC-FS FC-2 serial frame and transmits it into the Fibre Channel link.

As part of building the FC-FS FC-2 frame for the I/O request, the FICON channel in FICON native (FC) mode constructs the 24-bit FC port address of the destination N\_Port of the control unit, and the control unit image and device address within the physical CU.

### **FICON frame layout**

The FC-4 defines the mapping rules used by FICON. FC-SB-3 and FC-SB-4 are based on the FC-4 Information Unit constructs. Information associated with the execution of an I/O operation and the operation of a device is transferred between the channel and control unit as Information Units. IUs contain device-level commands, status, data, or control information or link-level control information. Figure 2-3 on page 27 illustrates the FICON frame layout.

The Fibre Channel (FC-2) sends the SB-3 and SB-4 IUs in FC-2 frames. The maximum amount of requester information that can fit into one FC frame is 2048 bytes (FICON does *not* use the optional header shown in the Payload in Figure 2-3 on page 27). As many frames as it takes to hold the IU data will be sent to transfer the IU. The receiving end node (FC adapter at the CU or another channel) rebuilds the incoming FC frames for a given IU into an inbound IU.

The FC header and its content determines how the frame is to be handled within the fabric. The frame contains two parts: the FC-FS FC-2 header part and the payload part (the FC-SB-4 or FC-SB-3 parts). The FC-2 header part of the frame holds the source FC port address (S\_ID) and the destination FC port address (D\_ID).

The FC-SB-3 or FC-SB-4 header part of the frame holds the channel image address, the Control Unit Image address, and the Device Unit Address (UA).

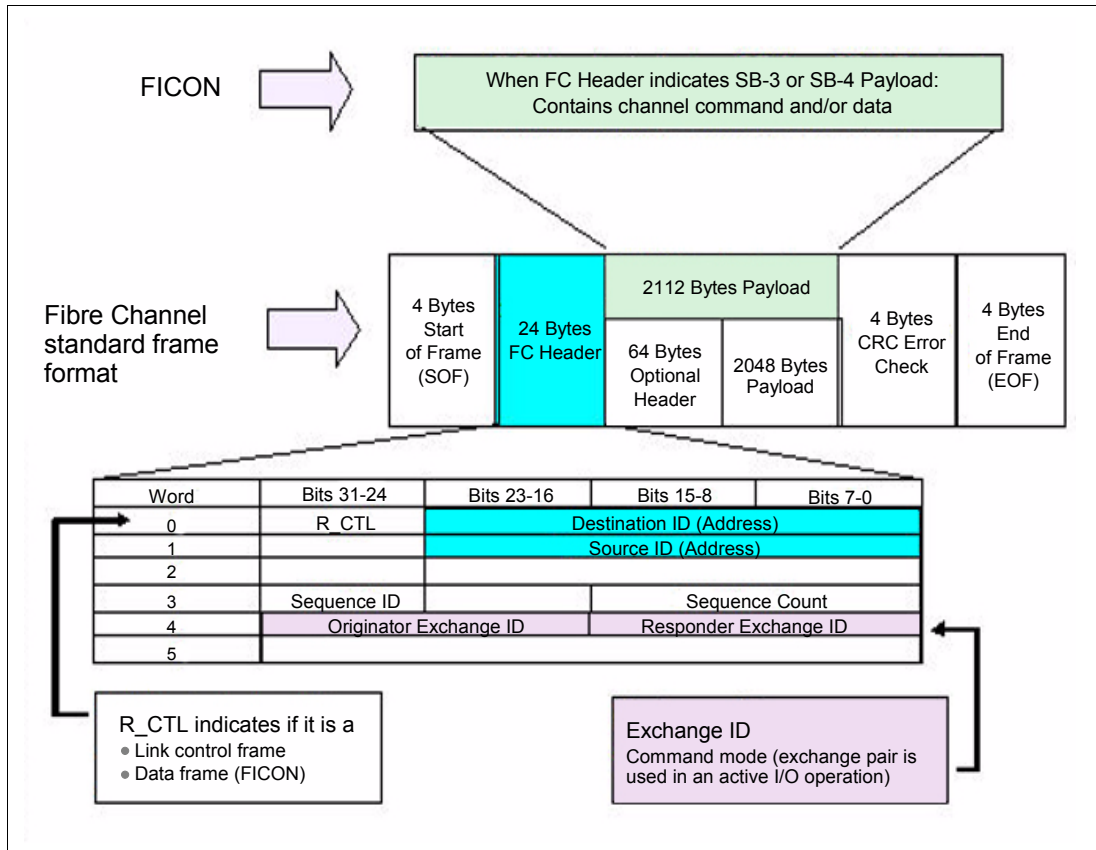


Figure 2-3 FICON frame layout

The FC header and its content determine how the frame is to be handled. The R\_CTL field indicates the type of frame it is, as follows:

- ▶ Link Control:
  - Acknowledge
  - Link Response
  - Link Command
- ▶ Data Frame:
  - Video\_Data
  - Link Data:
    - Basic Link Service (BLS)
    - Extended Link Service (ELS)
  - Device\_Data (Type - IP, IPI-3, SCSI, FCP, SB-3 or SB-4):
    - SB-3 and SB-4 is used for FICON

The FC header also contains the Exchange ID. In FICON transport mode, two exchange IDs are used and are called an *exchange pair*. Command mode uses one exchange pair for one I/O operation. An *open exchange* represents an I/O operation in progress over the channel.

A detailed description of the FC frame and FICON FC-SB-3 or SB-4 usage documentation can be obtained from the following web site:

<http://www.t11.org>

## 2.1.3 Command mode

The System z FICON architecture defines the protocol for CCW, data pipelining, and how FICON channels operate (known as command mode). Command mode uses the CCWs and data, as shown in Figure 2-4, to perform I/O operations.

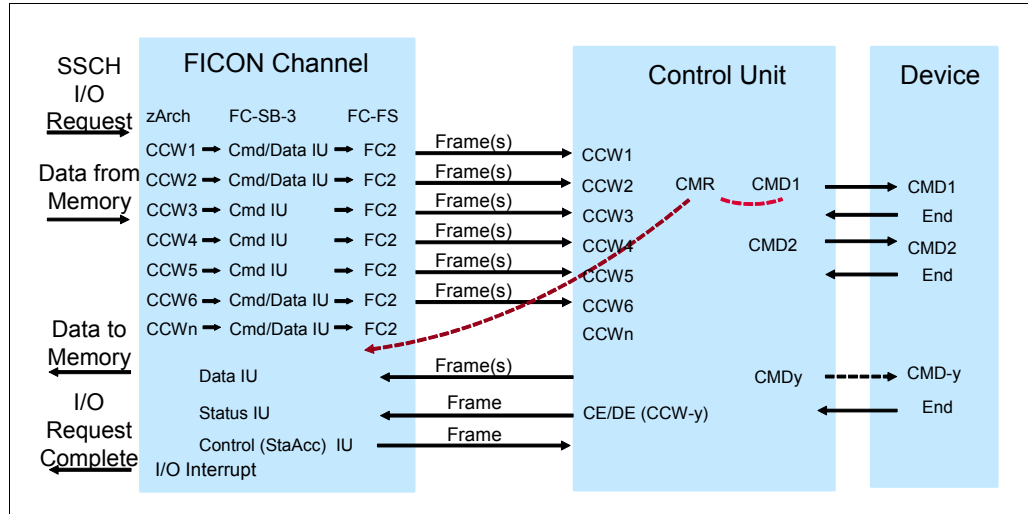


Figure 2-4 FICON command mode operation

A fundamental difference with ESCON is the CCW chaining capability of the FICON architecture. ESCON channel program operation requires a Channel End/Device End (CE/DE) after executing each CCW. FICON supports CCW chaining without requiring a CE/DE at the completion of each CCW operation.

The ESCON channel transfers the CCW to the control unit and waits for a CE/DE presented by the control unit after execution of the CCW by the device (CCW interlock). After receiving CE/DE for the previous CCW, the next CCW is transferred to the control unit for execution.

With a FICON channel, CCWs are transferred to the control unit without waiting for the first command response (CMR) from the control unit or for a CE/DE after each CCW execution. The device presents a logical end to the control unit after each CCW execution. After the last CCW of the CCW chain is executed by the CU/device, the control unit presents a CE/DE to the channel.

In addition, FICON channels can multiplex data transfer for several devices at the same time. This also allows workloads with low to moderate control unit cache hit ratios to achieve higher levels of activity rates per channel.

## 2.1.4 Transport mode

The way transport mode handles an I/O operation is significantly different from the CCW operation for command mode. While in command mode, each single CCW is sent to the control unit for execution. In transport mode, all commands are sent in a single frame to the control unit.

The zHPF channel program is divided into four parts, as shown in Figure 2-5 on page 29:

1. Transport Control Word (TCW): Used by the channel to locate the other parts of the channel program.

2. Transport Command Control Block (TCCB): Contains the commands and control data (for example, define extent parameters) used by the control unit.
3. Data: The buffers containing the data to be sent to the control unit (write) or the buffers to contain the data received from the control unit (read). The data buffers might be scattered in storage and addressed through a Transport Indirect Address List (TIDAL):
  - TIDAL is similar to an Indirect Address List (IDAL) or Modified Indirect Address List (MIDAL) for CCW programs
4. Transport Status Block (TSB): Contains status information for the I/O operation that is not presented in the Interruption Response Block (IRB). This information is for operating system use only.

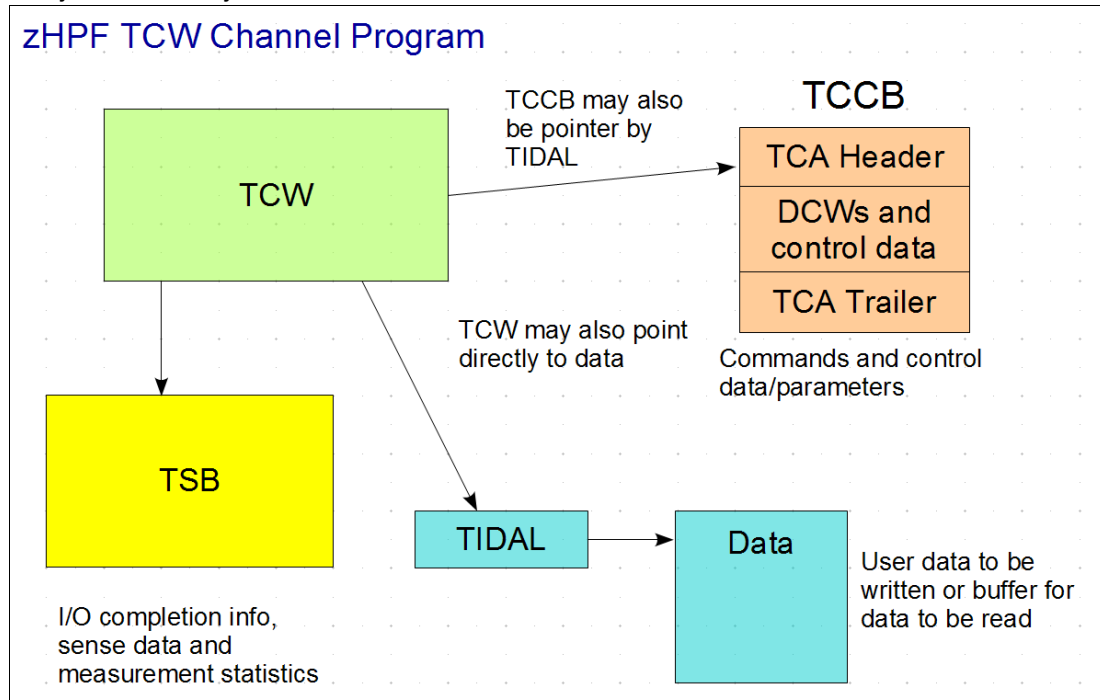


Figure 2-5 zHPF TCW Channel Program

The FICON channel uses transport mode if the control unit supports zHPF. When operating in transport mode, an I/O operation uses a transport control word to send commands and data, as illustrated by Figure 2-6 on page 30.

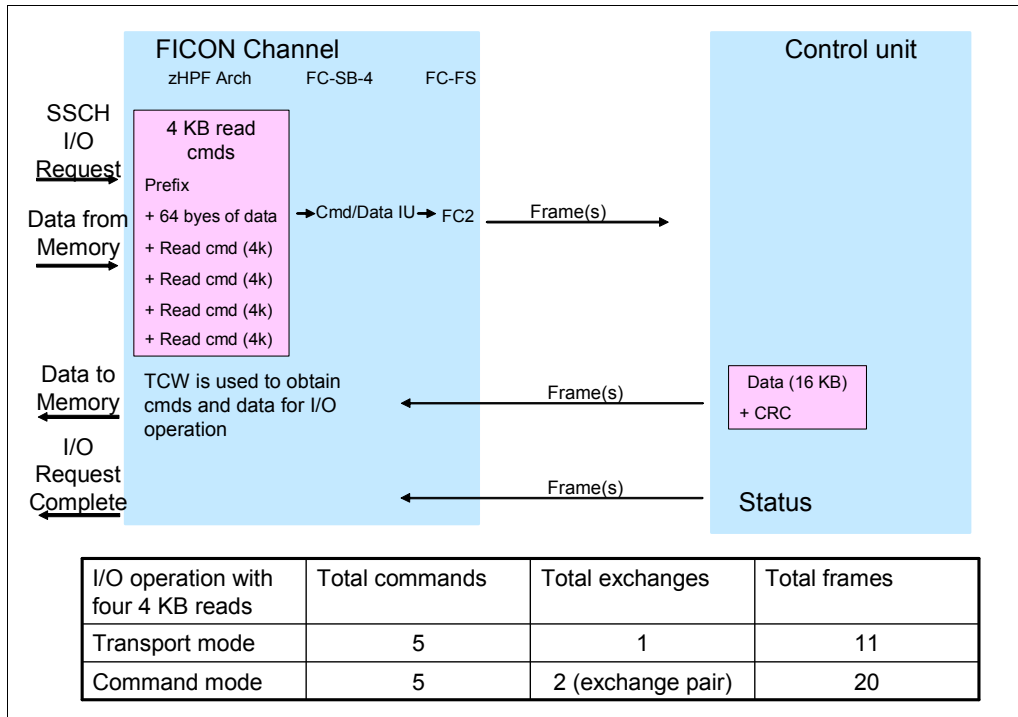


Figure 2-6 FICON transport mode operation

The FICON Express8S, FICON Express8, FICON Express4, and FICON Express2 features support transport mode. A parameter in the Operation Request Block (ORB) determines how the FICON channel will operate (command mode or transport mode). The mode used for an I/O operation also depends on the settings in the z/OS operating system.

The application or access method provides the channel program commands and parameters in the ORB. Bit 13 in word 1 of the ORB specifies how to handle the channel program in either command mode or transport mode.

### The zHPF support determination

The zHPF support is determined at link initialization. Figure 2-7 on page 31 shows the initialization steps that determine if transport mode is used.

During link initialization, both the channel and the control unit indicate whether they support zHPF. In the response to the Request Node ID (RNID) Extended Link Sequence (ELS), the Process Login (PRLI) support indicator is presented. In the RNID response, Process Login (PRLI) support is indicated in the node parameters, bit three:

- ▶ When bit three is one (1):
  - The node supports the FC-SB-4 process login using the PRLI extended link sequence (ELS):
    - FC-SB-4 protocol is used by zHPF
- ▶ When bit three is zero (0):
  - The node does not support the FC-SB-4 process login.

If PRLI is supported, the channel sends a PRLI ELS. The PRLI response then indicates that zHPF is supported by the control unit.

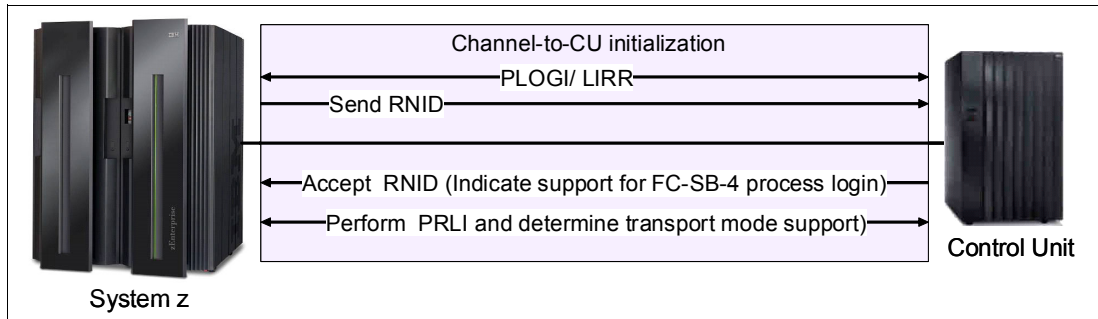


Figure 2-7 FICON sequence to determine zHPF support

A channel that supports the process login (PRLI) extended link service (ELS) and transport mode operations sends a request to each control unit in its configuration that also supports the process login ELS. The PRLI determines whether the control unit supports transport mode operations.

zHPF provides additional capabilities, such as being able to interrogate the CU before an actual missing interrupt occurs. Because zHPF uses TCWs (which do not time the channel operation in the same way as command mode), a mechanism must be used to prevent unnecessarily invoking the Missing Interrupt Handler (MIH) recovery actions. Transport mode provides an in-band method for z/OS to query the state of the I/O at the control unit without invoking error recovery or retry.

## 2.1.5 Missing Interrupt Handler

In channel operations, there is a Missing Interrupt Handler (MIH) time value set for various device classes. The MIH constantly polls and determines if there are any I/O operations that are not completing. For each of these I/O requests that are not completing, a recovery action is performed. The recovery action can be one or more of the following:

- ▶ Issue a message
- ▶ Generate a logrec record for diagnostic purposes
- ▶ Terminate the I/O request
- ▶ Requeue the I/O request

Normally when a channel does not get a CMR response from a started I/O operation within the MIH value, it invokes MIH recovery procedures, which can then terminate the I/O in progress.

The transport mode operation called *interrogate* queries the state of the I/O request from the CU. (The interrogate process does not affect the state of the primary operation and is sent before the MIH value is exceeded.) The CU provides an interrogate response IU that contains extended status that describes the state of the primary operation. z/OS can then decide whether recovery is needed or to reset MIH value for this I/O operation. The interrogate information appears in MIH messages and LOGREC records.

Interrogate is supported for device and driver specified MIH times, *not* I/O timing. IOS does not issue an interrogate if:

- ▶ MIH is disabled for the device or I/O request
- ▶ The MIH time is less than five seconds
- ▶ A Halt Subchannel (HSCH) or Clear Subchannel (CSCH) is active
- ▶ A hyperswap is active

- ▶ An I/O timeout is occurring during this or the next MIH interval
- ▶ The subchannel status is device active only (missing secondary status)

## 2.1.6 zHPF Temporary Logout

With this new function available using High Performance FICON for System z (zHPF) the control unit tells the z196 or z114 that the host must not start any new I/O until the specified temporary logout time expires. For DS8000 the temporary logout time is four seconds.

This process logout (PRLO) function is used during concurrent load or warmstart recovery scenarios.

For a code load, a PRLO (logout) is sent to the channel indicating that the channel must not send any new I/O for four seconds.

In an error recovery situation (warmstart), the control unit sends a PRLO (logout) to the channel with a time delay of one (1) second, indicating that the channel must not drive any new I/O for one second.

These times are fixed in the code and are not programmable. The channels specify the maximum time they can delay initiating I/O at process login (PRLI) time. All code levels that support zHPF utilize this feature.

## 2.1.7 Modified Indirect Data Address Word

On System z, the Modified Indirect Data Address Word (MIDAW) provides alternatives to using CCW data chaining in channel programs. The MIDAW facility was added to z/Architecture and can coexist with the current CCW IDAW facility.

MIDAW allows scattering of data in memory for non-contiguous real pages. (This is sometimes known as scatter-read or scatter-write.) Although the CCW IDAW function requires all but the first and last IDAW in a list to deal with complete 2 KB or 4 KB units of data, the MIDAW facility allows page boundary crossing on either 2 KB or 4 KB boundaries. This allows access to data buffers anywhere in a 64-bit buffer space. Figure 2-8 on page 33 illustrates an example of MIDAW usage.



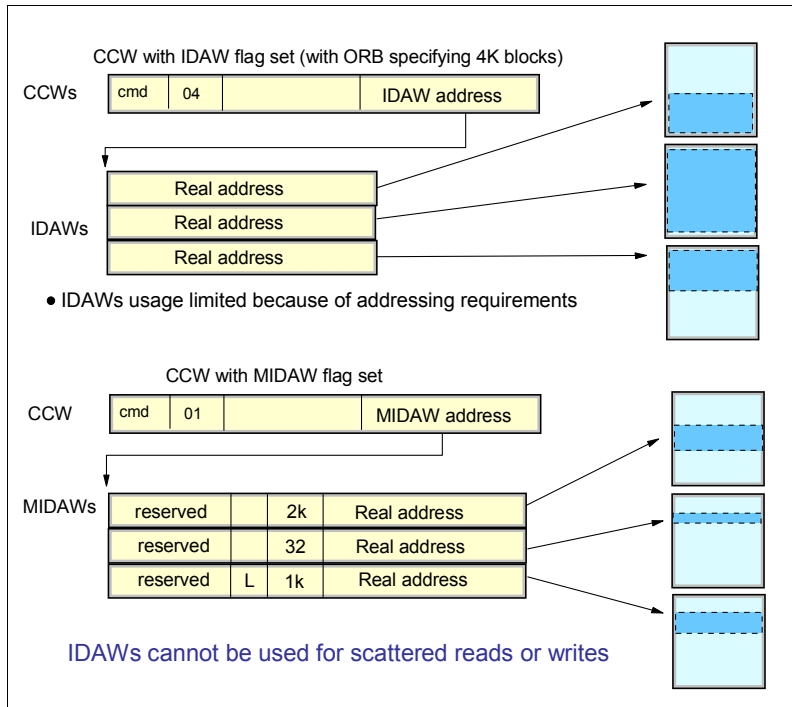


Figure 2-8 Command mode MIDAWs

The use of MIDAWs is indicated by the MIDAW bit (flag bit 7) in the CCW. The last MIDAW in the list has a *last* flag set, indicated by the L in Figure 2-8

MIDAWs provide significant performance benefits, particularly when processing extended format data sets with FICON channels.

## 2.1.8 Transport Indirect Data Address Word

When the System z channel is operating in transport mode, Transport Indirect Data Address Word (TIDAW) is used. TIDAWs and MIDAWs are similar in concept and provide the capability of scattered reads and writes.

TIDAWs are used when certain flag bits are set in the transport control word. Figure 2-9 illustrates an example of TIDAW usage.

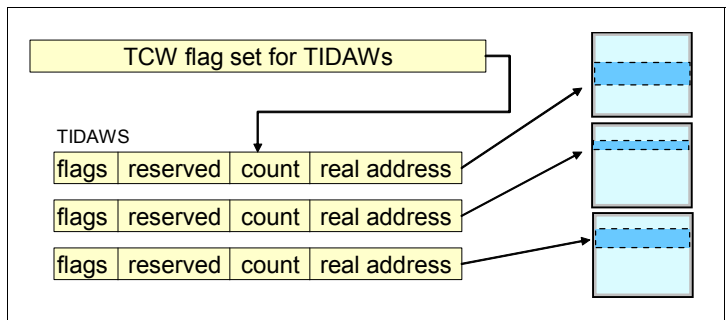


Figure 2-9 Transport mode TIDAWs

TIDAWs and MIDAWs are used with z/OS extended format data sets that use internal structures (usually not visible to the application program) that require scatter-read or scatter-write operation.

## 2.1.9 Open exchange

An open exchange is part of FICON (and FC) terminology. Many I/O operations can be in progress over FICON channels at any one time, for example, a disk I/O operation might temporarily disconnect from the channel while performing a seek operation or while waiting for a disk rotation. During this disconnect time, other I/O operations can be managed, as explained here:

- ▶ Command mode open exchanges

In command mode, the number of open exchanges is limited by the FICON Express feature. FICON Express8S, FICON Express8, FICON Express4, FICON Express2 allow up to 64 open exchanges. One open exchange (actually an exchange pair) in command mode is the same as one I/O operation in progress.

- ▶ Transport mode open exchanges

In transport mode, one exchange is sent from the channel to the control unit. Then the same exchange ID is sent back from the control unit to the channel to complete the I/O operation. The maximum number of simultaneous exchanges the channel can have open with a physical CU is 750 exchanges. The CU sets the maximum number of exchanges in the status area of the transport mode response IU (the default number is 64 and can be increased or decreased).

## 2.1.10 Buffer-to-buffer credit usage in FICON

Normally buffer credits are only an issue for extended distances. The buffer-to-buffer value determines the distance that the two nodes can be apart and still maintain the supported link data rate. All System z FICON channel features have enough buffer credits to support an unrepeated distance of 10 km.

**Note:** FICON channels experience minimal data rate droop at repeated distances up to 100 km (62 miles). While it is theoretically possible for FICON to maintain high bandwidth at distances greater than 100 km, these distances have not been System z qualified and are only achievable if enough buffer credits exist to support the link speed.

When FICON frames travel over the link, a flow control mechanism must be in place to ensure the transmitter does not overrun the receiver. FICON mainly delivers class 3 type frames that use buffer-to-buffer credits to manage the flow control

Ports have buffers for the temporary storage of frames. Each buffer can contain up to 2112 bytes of data (FICON does *not* use the optional header shown in the Payload in Figure 2-3 on page 27). Buffers are referred to as *credits*. A credit represents a receiving port's ability to accept one frame. At port initialization (login), buffer credit values are exchanged between two ports based on the number of buffers available for the ports. For example, the FICON Express8S and FICON Express8 features have 40 buffer credits, which is 80 KB of buffer at the receiver. This means that the FICON Express8S and FICON Express8 receivers can store up to 40 frames of data at any one time. The other FICON features support:

- ▶ FICON Express4 contains 212 buffer credits.
- ▶ FICON Express2 contains 107 buffer credits.
- ▶ FICON Express contains 64 buffer credits.

Buffer-to-buffer credit is on a FRAME basis, not a FRAME SIZE basis. A one-byte frame consumes one credit as does a two-byte frame.

The number of buffer credits required depends on the link data rate, the distance between two adjacent nodes, the average frame size, the maximum number of buffer credits supported by the FICON director or control unit, and the application and workload characteristics.

For a more complete explanation about how buffer credits work, refer to “Buffer credits” on page 51.

### 2.1.11 Extended distance FICON

Degradation of performance at extended distances can be avoided by implementing an enhancement to the industry standard FICON architecture. The enhancement is a protocol for *persistent* Information Unit (IU) pacing, and it can be used when the FC-SB-4 process login is supported. Control units that exploit the architecture can increase the pace count (the number of IUs allowed to be underway between the channel and the control unit). Extended distance FICON remembers the last pacing information and uses this information for subsequent operations, thus avoiding performance degradation at the start of a new operation.

The IU pacing protocol controls the number of IUs that can be in flight from a channel to a control unit. The control unit might increase the pacing count (the number of IUs allowed to be in flight from a channel to a control unit) in the first Command Response IU sent to the channel. When the enhancement for persistent IU pacing is not used, the increased pacing count is valid only for the remainder of the current outbound exchange. See Figure 2-10.

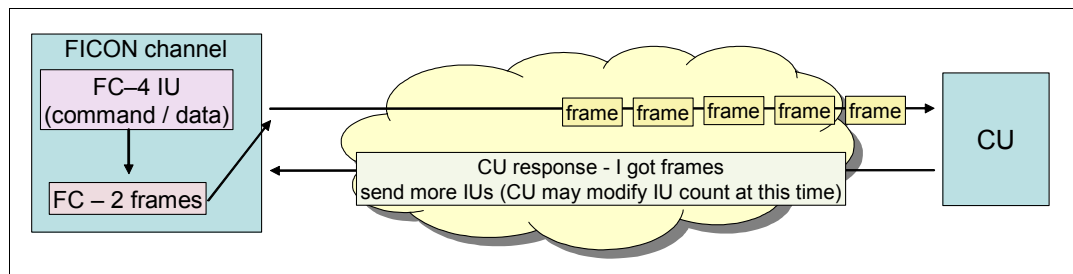


Figure 2-10 IU pacing for extended distance

At the start of every I/O, the IU count is reset to the default number. The IU pacing protocol, as defined, has the limitation that the first burst of IUs from the channel to the control unit can be no larger than a default value of 16. This setting causes a delay in the execution of channel programs with more than 16 IUs at long distances because a round trip to the control unit is required before the remainder of the IUs can be sent by the channel, upon the receipt of the first command response, as allowed by the increased pacing count.

A channel can operate in the default IU pacing mode or in the persistent IU pacing mode. During initialization, the specific node descriptor information is exchanged between the channel and control unit. This information includes SB-4 support and indicates whether the node supports concurrent enablement of the persistent IU pacing function:

- ▶ When a channel that supports concurrent enablement of the persistent IU pacing function receives a node descriptor from a control unit with bit 7 of byte 1 equal to one, the channel enables persistent IU pacing for all currently established logical paths with the control unit.
- ▶ When a control unit that supports concurrent enablement of the persistent IU pacing function receives a node descriptor from a channel with bit 7 of byte 1 equal to one, the

control unit enables persistent IU pacing for all currently established logical paths with the channel.

- ▶ For logical paths that are established subsequent to the processing of the node descriptor, the persistent IU pacing bit in the Establish Logical Path (ELP)/Logical Path Established (LPE) IU optional features field is used to enable or disable persistent IU pacing.

Persistent IU pacing is a method for allowing a channel and control unit supporting the FC-SB-4 process login to retain a pacing count that can be used at the start of execution of a channel program. This can improve performance of long I/O programs at higher link speeds and long distances by allowing the channel to send more IUs to the control unit and eliminating the delay of waiting for the first Command Response.

The channel retains the pacing count value, presented by the control unit in accordance with the standard, and uses that pacing count value as its new default pacing count for any new channel programs issued on the same logical path. Figure 2-11 illustrates an example of using persistent IU pacing for extended distance.

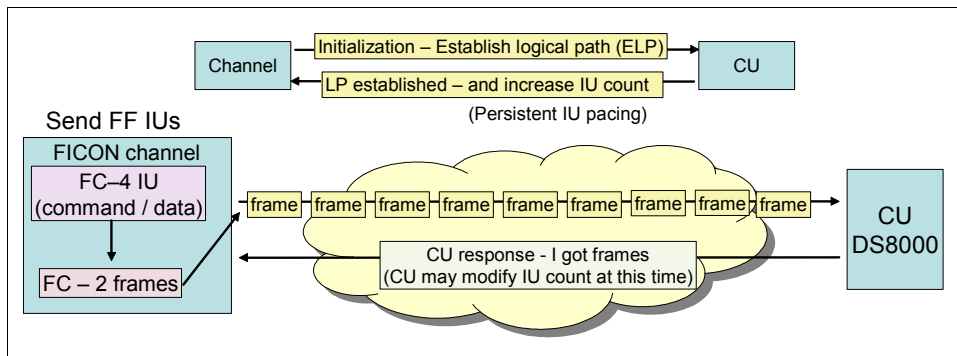


Figure 2-11 Persistent IU pacing for extended distance

Extended distance FICON is transparent to operating systems and applies to all FICON Express8S, FICON Express8, and FICON Express4 features carrying native FICON traffic (CHPID type FC).

To exploit extended distance FICON, the control unit must support the new IU pacing protocol. IBM System Storage® DS8000 with the appropriate licensed machine code level supports extended distance FICON on z196, z114, and z10 servers.

**Note:** The channel defaults to current pacing values when it operates with control units that cannot exploit extended distance FICON.

## 2.1.12 FICON link incident reporting

FICON link incident reporting allows an operating system image to register link incident reports without operator intervention.

To capture this data for link error analysis, z/OS can now register to receive FICON link incident records, referred to as Link Maintenance Information (LMI) records. z/OS displays this information that is on consoles and saves it in the system log and in LOGREC.

FICON link incident records consist of detailed node information that is gathered for particular link-degraded and link-failure incidents. The node descriptor information is gathered during the Fabric login (FLOGI) or Port login (PLOGI) when the logical path is established.

For additional information pertaining to LMI records, refer to *z/OS MVS Diagnosis: Reference*, GA22-7588.

## 2.2 System z FICON feature support

Table 2-1 lists the FICON features that are available on the z196, z114 and z10 servers, and the number of supported FICON channels for each. Some features are only available when carried forward on a server upgrade, and FICON Express8S features are only available on the zEnterprise z196 and z114 servers.

The channels residing on a single FICON Express8S, FICON Express8, FICON Express4, or FICON Express2 feature can be configured individually and can be defined in several channel modes (FC or FCP).

The channels residing on a FICON Express feature can be defined for FC, FCP, or FCV channel mode; however, this feature is only supported on a System z10 server if the feature was carried forward on an upgrade.

Table 2-1 System z server FICON feature support

| Channel feature                     | Feature codes | Maximum number of channels |                  |      |      | Channels per feature | Channel increments (orderable) |
|-------------------------------------|---------------|----------------------------|------------------|------|------|----------------------|--------------------------------|
|                                     |               | z10 EC                     | z10 BC           | z196 | z114 |                      |                                |
| FICON Express LX <sup>a</sup>       | 2319          | 120                        | 40               | n/a  | n/a  | 2                    | 2                              |
| FICON Express SX <sup>a</sup>       | 2320          | 120                        | 40               | n/a  | n/a  | 2                    | 2                              |
| FICON Express2 LX                   | 3319          | 336 <sup>a</sup>           | 112 <sup>a</sup> | n/a  | n/a  | 4                    | 4                              |
| FICON Express2 SX                   | 3320          | 336 <sup>a</sup>           | 112 <sup>a</sup> | n/a  | n/a  | 4                    | 4                              |
| FICON Express4 10KM LX <sup>b</sup> | 3321          | 336                        | 128              | 288  | 64   | 4                    | 4                              |
| FICON Express4 SX <sup>b</sup>      | 3322          | 336                        | 128              | 288  | 64   | 4                    | 4                              |
| FICON Express4-2C SX                | 3318          | n/a                        | 64               | n/a  | 32   | 2                    | 2                              |
| FICON Express4-2C 4KM LX            | 3323          | n/a                        | 64               | n/a  | 32   | 2                    | 2                              |
| FICON Express4 4KM LX <sup>b</sup>  | 3324          | 336                        | 128              | 288  | 64   | 4                    | 4                              |
| FICON Express8 10KM LX              | 3325          | 336                        | 128              | 288  | 64   | 4                    | 4                              |
| FICON Express8 SX                   | 3326          | 336                        | 128              | 288  | 64   | 4                    | 4                              |
| FICON Express8S 10 KM LX            | 0409          | n/a                        | n/a              | 320  | 128  | 2                    | 2                              |
| FICON Express8S SX                  | 0410          | n/a                        | n/a              | 320  | 128  | 2                    | 2                              |

a. Carry forward on an upgrade

b. Effective October 27, 2009, withdrawn from marketing for System z10 servers

### FICON Express8S

The FICON Express8S features reside exclusively in the PCIe I/O drawer available only on z196 and z114. Each of the two independent ports is capable of 8 gigabits per second (Gbps), 4 Gbps, or 2 Gbps depending on the capability of the attached device or switch. The link speed is auto-negotiated, point-to-point, and is transparent to users and applications. A link data rate of 1 Gbps is *not* supported.

All FICON Express8S features use Small Form Factor Pluggable (SFP) optics to permit each channel to be individually serviced in the event of a fiber optic module failure. The traffic on the other channels on the same feature can continue to flow if a channel requires servicing.

Each FICON Express8S feature occupies one I/O slot in the Peripheral Component Interconnect express (PCIe) I/O drawer and each has two ports, both supporting LC Duplex connectors. Each feature has one PCHID and one CHPID associated to each port. It allows the consolidation of existing FICON Express8, FICON Express4, FICON Express2, and FICON Express channels onto fewer FICON Express8S channels while maintaining and enhancing performance.

FICON Express8S offers performance improvements for zHPF and FCP to ensure that the z196 and z114 servers continue to allow your bandwidth to increase to meet the demands of your business applications. For more information about FICON channel performance, see the technical papers on the System z I/O connectivity web site at:

[http://www-03.ibm.com/systems/z/hardware/connectivity/ficon\\_performance.html](http://www-03.ibm.com/systems/z/hardware/connectivity/ficon_performance.html)

FICON Express8S CHPIDs can be defined as spanned channels and can be shared among logical partitions within and across CSSs. They also support cascading (the connection of two FICON directors using Inter-Switch Links or ISLs) to minimize the number of cross-site connections and help reduce implementation costs for disaster recovery (DR) sites, PPRC, IBM GDPS®, and remote copy.

The FICON Express8S features are designed for connectivity to servers, switches, Directors, disks, tapes, and printers, and they can be defined as:

- ▶ CHPID type FC:
  - FICON, High Performance FICON for System z (zHPF), and FICON Channel-to-Channel (FCTC) traffic and these protocols are supported simultaneously
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments
- ▶ CHPID type FCP:
  - Fibre Channel Protocol traffic for communication with SCSI devices
  - Supported in z/VM, z/VSE, and Linux on System z environments

The two types of FICON Express8S optical transceivers supported are the long wavelength (LX) and the short wavelength (SX). All ports on a FICON feature must be of the same type, either LX or SX.

### ***FICON Express8S 10KM LX***

Feature code 0409 has two ports per feature and supports unrepeated distances up to 10 km (6.2 miles) at 8 Gbps.

Each channel supports 9/125 micrometer single mode fiber optic cable terminated with an LC Duplex connector.

### ***FICON Express8S SX***

Feature code 0410 has two ports per feature and supports unrepeated distances up to 150 meters (492 feet) at 8 Gbps.

Each channel supports 50/125 or 62.5/125 micrometer multimode fiber optic cable terminated with an LC Duplex connector.

## **FICON Express8**

The FICON Express8 features are designed to deliver increased performance compared to the FICON Express4 and FICON Express2 features.

The FICON Express8 features have four independent channels, and each feature occupies a single I/O slot, utilizing one CHPID per channel. Each channel supports 8 Gbps, 4 Gbps, and 2 Gbps link data rates with auto-negotiation. A link data rate of 1 Gbps is *not* supported.

All FICON Express8 features use Small Form Factor Pluggable (SFP) optics to permit each channel to be individually serviced in the event of a fiber optic module failure. The traffic on the other channels on the same feature can continue to flow if a channel requires servicing.

The FICON Express8 features are ordered in 4-channel increments and are designed to be added concurrently. This concurrent update capability allows you to continue to run workloads through other channels while the FICON Express8 features are being added.

FICON Express8 CHPIDs can be defined as spanned channels and can be shared among logical partitions within and across CSSs. They also support cascading (the connection of two FICON directors in succession) to minimize the number of cross-site connections and help reduce implementation costs for disaster recovery (DR) sites, PPRC, GDPS, and remote copy.

FICON Express8 features are designed for connectivity to servers, switches, Directors, disks, tapes, and printers, and they can be defined as:

- ▶ CHPID type FC:
  - FICON, High Performance FICON for System z (zHPF), and FICON Channel-to-Channel (FCTC) traffic and these protocols are supported simultaneously
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments
- ▶ CHPID type FCP:
  - Fibre Channel Protocol traffic for communication with SCSI devices
  - Supported in z/VM, z/VSE, and Linux on System z environments

### ***FICON Express8 10KM LX***

Feature code 3325, with four ports per feature, is designed to support unrepeated distances up to 10 km (6.2 miles) at 8 Gbps.

Each channel supports 9/125 micrometer single mode fiber optic cable terminated with an LC Duplex connector.

### ***FICON Express8 SX***

Feature code 3326, with four ports per feature, is designed to support unrepeated distances up to 150 meters (492 feet) at 8 Gbps.

Each channel supports 50/125 micrometer multimode fiber optic cable or a 62.5/125 micrometer multimode fiber optic cable terminated with an LC Duplex connector.

## **FICON Express4**

The FICON Express4 features have four (or two, for the 2-port features) independent channels, and each feature occupies a single I/O slot, utilizing one CHPID per channel. Each channel supports 4Gbps, 2 Gbps, and 1Gbps link data rates with auto-negotiation.

All FICON Express4 features use Small Form Factor Pluggable (SFP) optics to permit each channel to be individually serviced in the event of a fiber optic module failure. The traffic on the other channels on the same feature can continue to flow if a channel requires servicing.

The FICON Express4 features are ordered in 4-channel (or 2-channel) increments and are designed to be added concurrently. This concurrent update capability allows you to continue to run workloads through other channels while the FICON Express4 features are being added.

FICON Express4 CHPIDs can be defined as a spanned channel and can be shared among logical partitions within and across CSSs. They also support cascading (the connection of two FICON directors in succession) to minimize the number of cross-site connections and help reduce implementation costs for disaster recovery (DR) sites, PPRC, GDPS, and remote copy.

The FICON Express4 features are designed for connectivity to servers, switches, Directors, disks, tapes, and printers, and they can be defined as:

- ▶ CHPID type FC:
  - FICON, High Performance FICON for System z (zHPF), and FICON Channel-to-Channel (FCTC) traffic and these protocols are supported simultaneously
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments
- ▶ CHPID type FCP:
  - Fibre Channel Protocol traffic for communication with SCSI devices
  - Supported in z/VM, z/VSE, and Linux on System z environments

### ***FICON Express4 10KM LX***

Feature code 3321 is designed to support unrepeated distances up to 10 km (6.2 miles) at 4 Gbps.

Each channel supports 9 micron single mode fiber optic cable terminated with an LC Duplex connector.

Multimode (62.5 or 50 micron) fiber cable can be used with the FICON Express4 10KM LX feature. The use of multimode cable types requires a mode conditioning patch (MCP) cable.

This feature must be used when the unrepeated distance between devices is greater than 4 km (2.5 miles) or the link loss budget between devices exceeds 2 dB. A 10KM LX transceiver is designed to interoperate with a 10KM LX transceiver. Interoperability of 10 km transceivers with 4 km transceivers is supported if the unrepeated distance does not exceed 4 km.

### ***FICON Express4 4KM LX***

Feature code 3324 is designed to support unrepeated distances up to 4 km (2.5 miles) at 4Gbps.

Interoperability of 4 km transceivers with 10 km(6.2 miles) transceivers is supported if the unrepeated distance does not exceed 4 km.

Each channel supports 9 micron single mode fiber optic cable terminated with an LC Duplex connector.

Multimode (62.5 or 50 micron) fiber cable can be used with the FICON Express4 4 KM LX feature. The use of multimode cable types requires a mode conditioning patch (MCP) cable.



***FICON Express4-2C 4KM LX*** (only available on z114 and z10 BC servers)

Feature code 3323, with two channels per feature, is designed to support unrepeated distances up to 4 km (2.5 miles) at 4 Gbps.

Interoperability of 4 km transceivers with 10 km (6.2 miles) transceivers is supported if the unrepeated distance does not exceed 4 km.

Each channel supports 9 micron single mode fiber optic cable terminated with an LC Duplex connector.

Multimode(62.5 or 50 micron) fiber cable can be used with the FICON Express4-2C 4 KM LX feature. The use of multimode cable types requires a mode conditioning patch (MCP) cable.

***FICON Express4 SX***

Feature code 3322 is designed to support unrepeated distances up to 270 meters (886 feet) at 4 Gbps.

Each channel supports 62.5 micron or 50 micron multimode fiber optic cable terminated with an LC Duplex connector.

***FICON Express4-2C SX***

This is only available on z114 and z10 BC servers.

Feature code 3318, with two channels per feature, is designed to support unrepeated distances up to 270 meters (886 feet) at 4 Gbps.

Each channel supports 62.5 micron or 50 micron multimode fiber optic cable terminated with an LC Duplex connector.

**Note:** It is intended that the z196 and z114 are the last System z servers to support FICON Express4 features. Clients must review the usage of their installed FICON Express4 channels and where possible migrate to FICON Express8S channels.

**FICON Express2**

The FICON Express2 SX and FICON Express2 LX features have four independent channels with each feature occupying a single I/O slot, utilizing one CHPID per channel and four CHPIDs per feature, while continuing to support 2Gbps and 1Gbps link data rates. The link speed is auto-negotiated.

The FICON Express2 SX and LX features are ordered in 4-channel increments and designed to be added concurrently. This concurrent update capability allows you to continue to run workloads through other channels while the FICON Express2 features are being added.

FICON Express2 CHPIDs can be defined as a spanned channel and can be shared among logical partitions within and across CSSs.

The FICON Express2 features are designed for connectivity to servers, switches, Directors, disks, tapes, and printers, and they can be defined as:

- ▶ CHPID type FC
  - Native FICON, High Performance FICON for System z (zHPF), and FICON Channel-to-Channel (FCTC) traffic
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments

- ▶ CHPID type FCP:
  - Fibre Channel Protocol traffic for communication with SCSI devices
  - Supported in z/VM, z/VSE, and Linux on System z environments

**Note:** FICON Express2 is supported on z10 servers only if carried forward on an upgrade. It is *not* available on z196 or z114.

### ***FICON Express2 LX***

Feature code 3319 is designed to support unrepeated distances up to 10 km (6.2 miles) at 2 Gbps.

Each channel supports 9 micron single mode fiber optic cable terminated with an LC Duplex connector.

Multimode (62.5 or 50 micron) fiber cable can be used with the FICON Express2 LX feature. The use of multimode cable types requires a mode conditioning patch (MCP) cable.

### ***FICON Express2 SX***

Feature code 3320 is designed to support unrepeated distances up to 500 meters (1640 feet) at 2 Gbps.

Each channel supports 62.5 micron or 50 micron multimode fiber optic cable terminated with an LC Duplex connector.

### **FICON Express**

The two channels residing on a single FICON Express feature occupy one I/O slot in the System z I/O cage. Each channel can be configured individually and support 1 Gbps link speed.

The FICON Express features are designed for connectivity to servers, switches, Directors, disks, tapes, and printers, and they can be defined as:

- ▶ CHPID type FC:
  - Native FICON and FICON Channel-to-Channel (FCTC) traffic
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments
- ▶ CHPID type FCP:
  - Fibre Channel Protocol traffic for communication with SCSI devices
  - Supported in z/VM, z/VSE, and Linux on System z environments
- ▶ CHPID type FCV:
  - Supported in the z/OS, z/VM, z/VSE, z/TPF, and Linux on System z environments

**Note:** FICON Express is supported on z10 servers only if carried forward on an upgrade. It is *not* available on z196 or z114.

### ***FICON Express LX***

Feature code 2319 is designed to support unrepeated distances up to 10 km (6.2 miles) at 1 Gbps.

Each channel supports 9 micron single mode fiber optic cable terminated with an LC Duplex connector.

Multimode (62.5 or 50 micron) fiber cable can be used with the FICON Express LX feature. The use of multimode cable types requires a mode conditioning patch (MCP) cable.

## FICON Express SX

Feature code 2320 is designed to support unrepeated distances up to 860 meters (2822 feet) at 1 Gbps.

Each channel supports 62.5 micron or 50 micron multimode fiber optic cable terminated with an LC Duplex connector.

Table 2-2 summarizes the available FICON feature codes and their respective specifications.

### Notes:

- ▶ Mode Conditioning Patch (MCP) cables can be used with FICON features that can operate at a link data rate of 1 Gbps (100 MBps) only.
- ▶ FICON Express8S and FICON Express8 features do not support the attachment of MCP cables.

Table 2-2 FICON channel specifications

| Feature code | Channel feature          | Connector type | Cable type   | Bit rate                     | Server                                     |
|--------------|--------------------------|----------------|--|------------------------------|--|
| 2319         | FICON Express LX         | LC Duplex      | SM 9 $\mu$ m   | 1 or 2 Gbps <sup>a</sup>     | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
|              |                          |                | with MCP <sup>c</sup> :<br>MM 50 $\mu$ m or<br>MM 62.5 $\mu$ m | 1 Gbps                       | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
| 2320         | FICON Express SX         | LC Duplex      | MM 62.5 $\mu$ m  | 1 or 2 Gbps <sup>a</sup>     | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
|              |                          |                | MM 50 $\mu$ m  |                              |  |
| 3319         | FICON Express2 LX        | LC Duplex      | SM 9 $\mu$ m   | 1 or 2 Gbps <sup>a</sup>     | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
|              |                          |                | with MCP <sup>c</sup> :<br>MM 50 $\mu$ m or<br>MM 62.5 $\mu$ m | 1 Gbps                       | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
| 3320         | FICON Express2 SX        | LC Duplex      | MM 62.5 $\mu$ m  | 1 or 2 Gbps <sup>a</sup>     | z10 EC <sup>b</sup><br>z10 BC <sup>b</sup> |
|              |                          |                | MM 50 $\mu$ m  |                              |  |
| 3318         | FICON Express4-2C SX     | LC Duplex      | MM 62.5 $\mu$ m  | 1, 2, or 4 Gbps <sup>a</sup> | z114<br>z10 BC                             |
|              |                          |                | MM 50 $\mu$ m  |                              |  |
| 3323         | FICON Express4-2C 4KM LX | LC Duplex      | SM 9 $\mu$ m   | 1, 2, or 4 Gbps <sup>a</sup> | z114<br>z10 BC                             |
|              |                          |                | with MCP <sup>c</sup> :<br>MM 50 $\mu$ m or<br>MM 62.5 $\mu$ m |                              |  |
| 3325         | FICON Express8 10KM LX   | LC Duplex      | SM 9 $\mu$ m   | 2, 4, or 8 Gbps <sup>d</sup> | z196<br>z114<br>z10 BC/EC                  |
| 3326         | FICON Express8 SX        | LC Duplex      | MM 62.5 $\mu$ m  | 2, 4, or 8 Gbps <sup>d</sup> | z196<br>z114<br>z10 BC/EC                  |
|              |                          |                | MM 50 $\mu$ m  |                              |  |

| Feature code | Channel feature            | Connector type | Cable type | Bit rate                        | Server       |
|--------------|----------------------------|----------------|------------|---------------------------------|--------------|
| 0409         | FICON Express8S<br>10KM LX | LC Duplex      | SM 9 μm    | 2, 4, or 8<br>Gbps <sup>d</sup> | z196<br>z114 |
| 0410         | FICON Express8S SX         | LC Duplex      | MM 62.5 μm | 2, 4, or 8<br>Gbps <sup>d</sup> | z196<br>z114 |
|              |                            |                | MM 50 μm   |                                 |              |

- a. Supports auto-negotiate with neighbor node
- b. Only supported when carried forward on an upgrade
- c. Mode conditioning patch cables might be used
- d. Supports auto-negotiate with neighbor node for link data rates 8, 4, and 2 Gbps only

**Note:** IBM does not support a mix of 50 μm and 62.5 μm fiber optic cabling in the same physical link.

Refer to Table 4-1 on page 85 for the allowable maximum distances and link loss budgets based on the supported fiber optic cable types and link data rates.



## FICON Director technical description

We provide a technical description of the common components and functions of a FICON Director that are needed to support FICON environments. We also offer information regarding FICON Directors that are qualified by IBM for System z use.

**Terminology:** Often the terms Director and switch are used interchangeably.

However, compared to a switch, a Director typically means more fully redundant hardware resulting in higher availability, larger port capacity, and additional capabilities.

## 3.1 The role of the FICON Director

The FICON Director provides a variety of features and functions. Its primary function is to allow the dynamic connections between System z servers, control units, or another FICON Director.

The FICON Director provides the ability to logically connect multiple control units to a set of channels from one or more System z servers, which in turn minimizes the cabling infrastructure and maximizes the connectivity. One might say that the FICON Director is at the heart of I/O connectivity for System z servers.

Figure 3-1 shows a conceptual view of frame processing in a FICON Director with multiple FICON channels and control units.

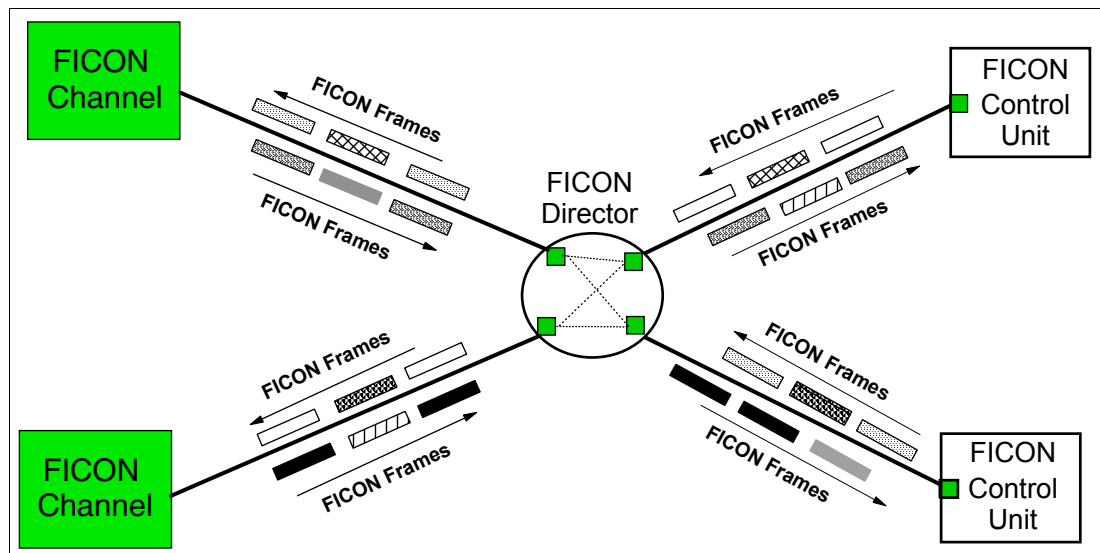


Figure 3-1 FICON Director - frame processing

The FICON channel supports multiple concurrent I/O connections. Each concurrent I/O operation can be to the same FICON control unit (but to different devices/control unit images), or to different FICON control units.

A FICON channel uses the Fibre Channel communication infrastructure to transfer channel programs and data through its FICON features to another FICON-capable node, such as a storage device, printer, or System z server (channel-to-channel). See Chapter 2, “System z FICON technical description” on page 21 for more details.

A FICON channel, in conjunction with the FICON Director, can operate in two topologies:

- ▶ Switched (through a single FICON Director to FICON-capable control units)
- ▶ Cascaded (through two FICON Directors to FICON-capable control units)

### 3.1.1 Switched configuration

In a FICON switched connection, two Fibre Channel (FC) links are needed in the channel-control unit path:

- ▶ Between the System z (N\_Port) and the FICON Director (F\_Port)
- ▶ Between the FICON Director (F\_Port) and the control unit (N\_Port)

The FICON channel determines whether the associated link is in a point-to-point or switched topology. It does so by logging into the fabric using fabric login (FLOGI ELS), and checking the accept response to the fabric login (ACC ELS). The FLOGI-ACC (accept) response indicates if the channel (N\_Port) is connected to another N\_Port (point-to-point) or an F\_Port (fabric port). Figure 3-2 shows an example of a switched topology.

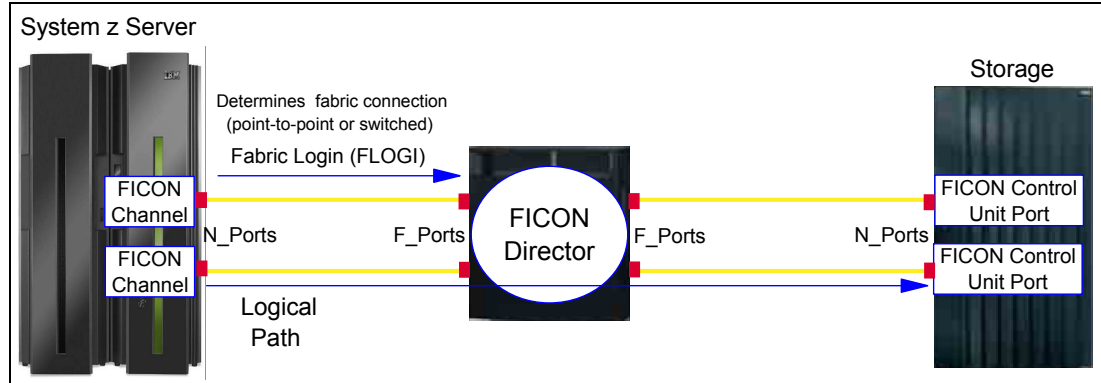


Figure 3-2 Switched configuration

Multiple channel images and multiple control unit images can share the resources of the FC link and the FICON Director, such that multiplexed I/O operations can be performed.

Channels and control unit links can be attached to the FICON Director in any combination, depending on configuration requirements and available resources in the Director.

Sharing a control unit through a FICON Director means that communication from a number of System z channels to the control unit can take place over one Director. A FICON channel can communicate with a number of FICON control units on different ports as well.

The communication path between a channel and a control unit is composed of two different parts: the physical channel path and the logical path.

In a FICON-switched topology (with a single Director), the physical paths are the FC links, or an interconnection of two FC links through a FICON Director, that provide the physical transmission path between a channel and a control unit.

A FICON (FC-SB-4) logical path is the relationship established between a channel image and a control unit image for communication during execution of an I/O operation and presentation of status.

### 3.1.2 Cascaded configuration

In a cascaded configuration, three Fibre Channel (FC) links are needed in the channel-control unit path:

- ▶ Between the System z (N\_Port) and the FICON Director (F\_Port)
- ▶ Between the two FICON Directors (E\_Ports)
- ▶ Between the FICON Director (F\_Port) and the control unit (N\_Port)

With this configuration, the connection between sites can consist of multiple FC links, also known as Inter-Switch Links (ISLs). An example of a cascaded topology is shown in Figure 3-3 on page 48.

Multiple channel images and multiple control unit images can share resources of the FC link and FICON Directors, such that multiplexed I/O operations can be performed.

Channels and control unit FC links can be attached to the FICON Directors in any combination, depending on configuration requirements and available Director ports.

Sharing a control unit through a FICON Director means that communication from a number of channels to the control unit can take place either over one Director-to-CU link (in the case where a control unit has only one FC link to the FICON Director), or over multiple link interfaces (in the case where a control unit has more than one FC link to the FICON Director).

A FICON channel can also communicate with a number of FICON control units on different ports in the second FICON Director.

The communication path between a channel and a control unit is composed of two separate parts: the physical channel path and the logical path.

In a cascaded topology, the physical paths are the FC links, interconnected by the Directors, that provide the physical transmission path between a FICON channel and a control unit. Figure 3-3 illustrates the configuration.

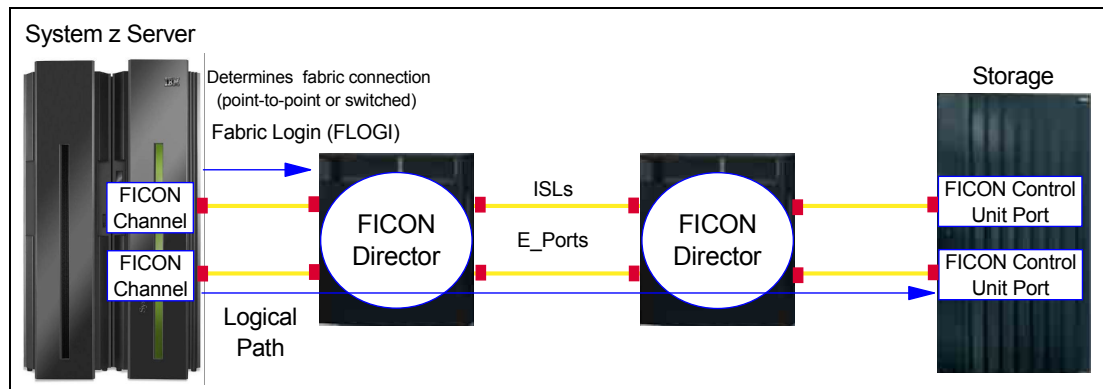


Figure 3-3 Cascaded configuration

A FICON (FC-SB-4) logical path is the relationship established between a FICON channel image and a control unit image for communication during execution of an I/O operation and presentation of status.

### 3.1.3 Basic components of a FICON Director

A FICON Director is an expandable and scalable Fibre Channel platform designed to support FICON attachment of System z servers and storage products. It typically features an N+1 redundancy scheme for all of its components. This redundancy results in high availability. Even in the case of failure of any single component, the FICON Director remains operational.

Most activities can be done non-disruptively, such as loading and activating firmware and replacing failed parts.

A FICON Director consists of a chassis that usually holds the following types of components:

- ▶ Backplane module
- ▶ Power Supply assembly
- ▶ Fan module assembly
- ▶ Control Processor module
- ▶ I/O module
- ▶ Switching module



## Backplane module

The backplane provides the connectivity within the chassis and between all system (I/O, switch, and control) modules, including the power supply assemblies and fan assemblies. The backplane gives you the capability to increase the number of ports by adding I/O modules.

### ***Oversubscription***

The term *oversubscription* implies that the port might not be able to drive at full speed due to limited backplane capacity. Oversubscription can only occur when simultaneous activities are present and backplane capacity is depleted, which usually applies to I/O modules with high port counts (for example, 48 ports operating at 8 Gbps).

## Power Supply assembly

Usually two power assemblies are installed in a chassis, providing full redundancy (and in some cases, load sharing). Each power supply assembly is a removable, hot-swappable unit.

## Fan module assembly

Fan module assemblies provide the cooling capability for the FICON Director. In the event of a single fan failure, the remaining fans can increase their speeds to a higher operating rpm to compensate for the single fan failure. The fan module assembly is hot-swappable.

## Control Processor module

Control Processor module contains microprocessor and memory, where the firmware resides and is responsible for Ficon Director manageability. The base configuration always contains two Control Processor modules for redundancy reasons. If the active switch module fails or is uninstalled, the standby Control Processor module automatically becomes the new active Control Processor module. Failover occurs as soon as the active Control Processor module is detected to be faulty or uninstalled.

The Control Processor module performs the following control functions:

- ▶ Switch initialization
- ▶ High availability
- ▶ Name server
- ▶ Simple Network Management Protocol (SNMP)
- ▶ Zoning Data Base
- ▶ CUP - In-band management of FICON Director

Typically, firmware upgrades can be performed concurrent to FICON Director operation.

## I/O module

The main function of the I/O module is to provide the physical connectivity between FICON channels and control units that are being connected to the Director.

Due to the usage of Small Form Factor Pluggable (SFP) ports instead of fixed optical port cards, port upgrades are simple and nondisruptive to other ports on the I/O module. Typically, any mix of short wavelength and long wavelength ports is allowed with any I/O module.

## Switching module

The switching module on the FICON Director contains the logic that provides overall switching in the Director. The switching module is responsible for core routing of frame within the FICON Director, although a frame can be switched to a destination FICON port without crossing the switching module, as implemented for some vendors. The base configuration always contains two switching modules.

### 3.1.4 Basic functions of a FICON Director

We discuss the basic functions of a FICON Director that are needed to support a FICON environment. Some common FICON Director terms are described too.

#### Domain ID or switch address

The Domain ID defines the FICON Director address (switch address in System z environment). It is specified in the FICON Director as a decimal number (must be converted to a hexadecimal value for use in the System z server). The switch address is a hexadecimal number used in defining the Ficon Director in the System z environment. The Domain ID and the switch address must be the same when referring to a Director in a System z environment.

The Domain ID and switch address must be unique. The switch address is used in the HCD, HCM, or IOCP definitions, and it can be any hex value between x'00' and x'FF'. The valid Domain ID range for the FICON Director is vendor-dependent, usually x'01' and x'EF'. When defining the switch addresses in HCD, HCM, or IOCP, ensure that you use values within the FICON Director's range.

The switch ID must be assigned by the user, and it must be unique within the scope of the definitions (HCD, HCM, or IOCP). The switch ID in the CHPID statement is basically used as an identifier or label. Although the switch ID can be different than the switch address or Domain ID, we recommend that you use the same value as the switch address and the Domain ID when referring to a FICON Director.

The Domain ID is assigned by the manufacturer, and it can be customized to another value. It must be unique within the fabric and insistent. *Insistent* means that the FICON Director uses the same Domain ID every time, even if it is rebooted. Insistent Domain IDs are required in a cascaded environment. There is no need for a Domain ID to ever change in a FICON environment, and fabrics come up faster after recovering from a failure if the Domain ID is insistent.

#### Port address or link address

The port address is a hexadecimal number between x'00' and x'FF' that is used in HCD, HCM, or IOCP to define a port on the FICON Director to which the device is connected, which is known as a one-byte link address. Some FICON Directors can modify the port address to any hexadecimal number that is required.

The port address is also needed in cascaded configurations, where it is part of the two-byte link address defined in the HCD, HCM, or IOCP. The two-byte link address consists of the Domain ID and the port address to which the device is connected.

#### Port number

The port number is typically used in open systems environments to refer to a port on a Director. In some cases, port numbers and port addresses are the same. However, in many cases they are contrasting, depending on the Director port characteristics. Port addressing is platform-specific and vendor-specific.

#### Inter-Switch Link

Inter-Switch Links (ISLs) are the fiber optic cables that connect two Directors using expansion ports (E\_Ports). ISLs carry frames originating from N\_Ports and frames that are generated within the fabric for management purposes.

Multiple ISLs are usually installed between Directors to increase bandwidth and provide higher availability.

## Fabric Shortest Path First

Fabric Shortest Path First (FSPF) is a Fibre Channel routing protocol that applies a weighting factor to paths based on link speed and status. FSPF associates a cost for each active link and computes paths from a Director to all the other Directors in the fabric. Paths are computed by adding the cost of all the links traversed by the path. Then a path is chosen with the lowest cost.

**Note:** FSPF only provides basic utilization of ISLs with the same cost. It does not ensure that all ISLs are equally utilized, and it does not inhibit congestion even if the other ISLs between Directors are not fully utilized.

## Control Unit Port

The Control Unit Port (CUP) protocol is used by System z to provide in-band management for FICON Directors. It allows z/OS to manage a FICON Director with the same level of control and security as for an ESCON Director. Host communication is through IBM Tivoli Systems Automation for z/OS (SA for z/OS), and includes control functions, such as blocking and unblocking of ports and monitoring and error reporting functions.

Although the FICON Director is transparent to the operating system in the path to a FICON control unit or device during the execution of an I/O operation, the CUP can provide:

- ▶ I/O management of the FICON Director
- ▶ FICON Director performance monitoring
- ▶ Error detection of the FICON Director

## Port swapping

Port swapping refers to the capability of a FICON Director to redirect traffic on a failed F\_Port to a working F\_Port without requiring a change in the System z I/O configuration.

In a Fibre Channel fabric, a failing port typically results in a cable being reconnected to another available port, followed by automatic discovery of the device through the name server. However, in a FICON environment, the control unit link address is defined in the channel configuration file (IOCP) of the System z server. Therefore, the FICON Director must ensure that the N\_Port address for the control unit remains the same even after the cable is reconnected to a different switch port.

After port swapping, the address assigned to the N\_Port performing fabric login (FLOGI) by the alternate port must be the same as the one that would have been assigned by the original port. The original port will assume the port address associated with the alternate port.

## Buffer credits

Buffer credits are used by the Fibre Channel architecture as a flow control mechanism to represent the number of frames a port can store.

Buffer credits are mainly a concern when dealing with extended distances and higher speeds. However, a poorly designed configuration can also consume available buffer credits and have a negative impact on performance.

Buffer credit flow control occurs at the link level (that is, between an N\_Port and an F\_Port or between two E\_Ports), and is used by Class 2 and Class 3 traffic. It relies on the receiver-ready (R\_RDY) signal to replenish credits. The total number of available buffer credits (BCs) that a receiver has is determined during the link initialization process (see Figure 3-4 on page 52).

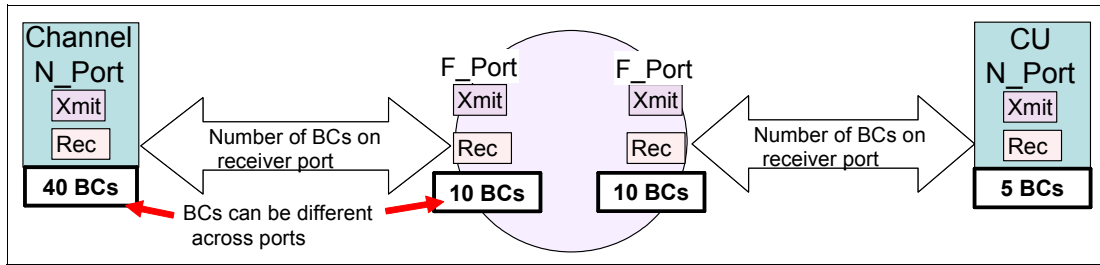


Figure 3-4 Buffer credit determination

The sender monitors the receiver's ability to accept frames by managing a buffer credit count value of the receiver port. The buffer credit count value is decremented by the transmitter when a frame is sent, as shown in Figure 3-5.

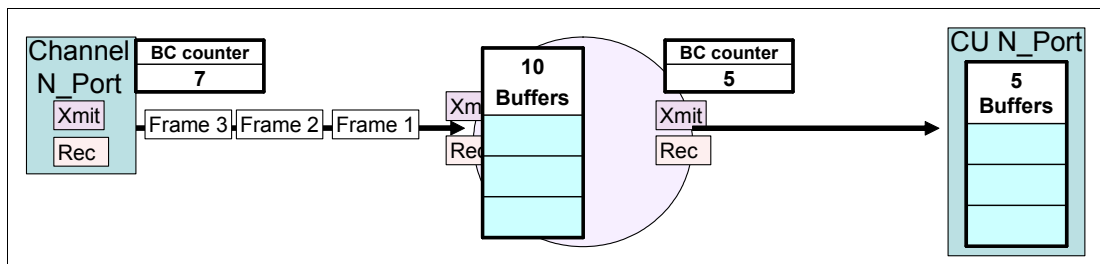


Figure 3-5 Buffer credit counters

An R\_RDY signal is returned for every received frame when a receive port buffer is made available. The buffer credit count value is incremented when an acknowledgment (R\_RDY) is received, as shown in Figure 3-6.

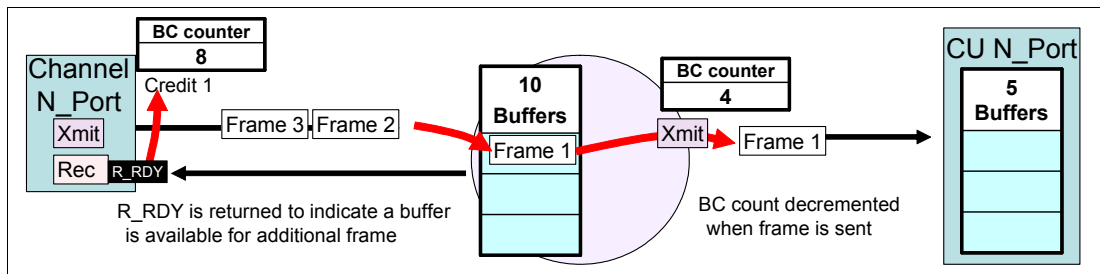


Figure 3-6 Buffer credit counters - R\_RDY

If the buffer credit count is zero (0), frame transmission to the associated port is temporarily suspended. This flow control mechanism prevents data overruns.

### Buffer credit calculation based on full frame size

To calculate the number of buffer credits needed to fully utilize an FC link with an average FICON frame size of 2048 bytes, use the following criteria:

- ▶ 1 buffer credit per kilometer (km) at 2 Gbps
- ▶ 2 buffer credits per kilometer (km) at 4 Gbps
- ▶ 4 buffer credits per kilometer (km) at 8 Gbps
- ▶ 8 buffer credits per kilometer (km) at 16 Gbps

In this case, a 10 km FC link running at 16 Gbps needs 40 buffer credits.

$$8 \text{ (per km at Gbps)} \times 10 \text{ km} = 80 \text{ buffer credits}$$

It is important to note that not all applications utilize the same frame size. FICON frame size information can be tracked using the FICON Director Activity Report in IBM RMF™. For RMF to receive the information, the CUP must be configured and the FDR parameter in SYS1.PARMLIB must be set to Yes.

The FICON maximum frame size of 2048 bytes can be achieved by using zHPF with large sequential reads/writes or with a large block size, such as 4 x 4 KB or 12 x 4 KB blocks per I/O.

### **High integrity fabric support**

Another important benefit of FICON support of cascaded Directors is its ability to provide high integrity data paths. The high integrity function is an integral component of the FICON architecture when configuring FICON channel paths through a cascaded fabric.

End-to-end data integrity is designed to be maintained through the cascaded environment. High integrity helps to ensure that any changes to the data streams are always detected, and that the data frames are delivered to the correct endpoint (FICON channel or control unit port). For FICON channels, Cyclical Redundancy Checking (CRC) and Longitudinal Redundancy Checking (LRC) are bit patterns added to the data streams to allow the detection of any bit changes.

High integrity features are used within the FICON channel and the cascaded Directors to ensure the detection and reporting of any miscabling actions occurring during operational use, which will cause data frames to be delivered to the wrong endpoint.

A FICON channel requires that the cascaded Directors support high integrity. During initialization, the FICON channel queries the fabric to determine that it supports high integrity. If it does, then the channel completes the initialization process, allowing the channel to operate with the fabric.

After a FICON switched fabric is customized to support FICON cascaded Directors, and the required World Wide Node Name (WWNN) and Domain IDs are added in the fabric membership list, the Director checks that its Inter-Switch Links (ISLs) are attached to the correct Director before they are made operational. If an accidental cable swap occurs, the Director invokes logical path testing, reporting, isolation, and recovery. The high integrity fabric feature for cascaded FICON Directors protects against miscabling and misdirecting of data streams, as shown in Figure 3-7 on page 54.

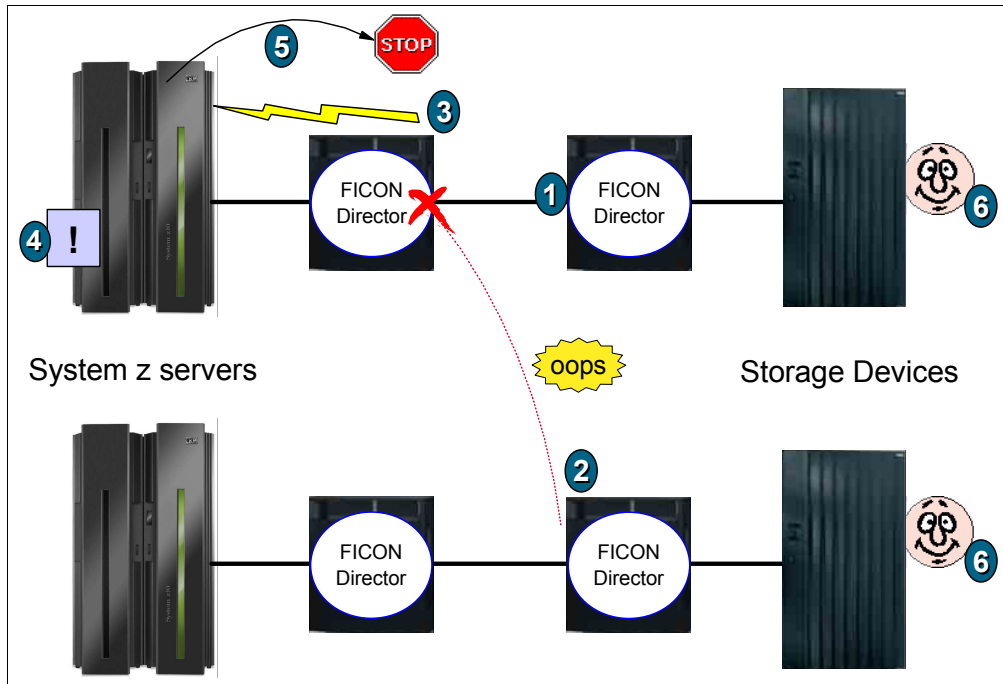


Figure 3-7 High integrity fabric feature

The checking process proceeds in this way:

1. Channel initialization completes.
2. Some later time, miscabling occurs (for example, cables are swapped at a patch panel).
3. The Director port enters “invalid attachment state” and notifies state change back to System z.
4. The System z server invokes the channel logical path testing, reporting, and isolation, and error recovery.
5. Any I/O requests to the invalid route are discarded until the error is corrected.
6. Data is protected.

### Fabric binding

Fabric binding is a security feature that enables explicit control over which FICON Directors can be interconnected by preventing non-authorized Directors from merging, either accidentally or intentionally.

This is done by manually defining the authorized Director in a fabric binding database. The fabric binding database contains the World Wide Node Name (WWNN) and Domain ID of the connecting FICON Director. The FICON Director that is allowed to connect must be added to the fabric binding database of the other FICON Director.

**Note:** Activating fabric binding is a prerequisite for a cascaded topology.

### Insistent Domain IDs

Insistent Domain ID support restricts a switch address (Domain ID) from being automatically changed when a duplicate switch address is added to the fabric. Overt operator action is required to change a switch address.

Insistent Domain ID support also prohibits the use of dynamic Domain IDs to ensure that predictable Domain IDs are enforced within the fabric. For example, if a Director has this feature enabled, and a new Director is connected to it (through an ISL) without the Domain ID feature enabled, the new Director is segmented into a separate fabric and user data will not flow.

## Zoning

Zoning is a method used in the FICON Director to enable or disable communication between different attached devices. A *zone* consists of a group of ports or WWNs. Connectivity is permitted only between connections to the devices that are in the same zone.

There are two types of zoning:

- ▶ Name zoning, which permits connectivity between attached nodes based on WWN
- ▶ Port zoning, which restricts port connectivity based on port number

WWN zoning is typically used for open systems connections. Port zones are used for FICON.

Nodes that are not already defined to a zone are automatically put into the default zone. Conversely, default zone members are automatically removed whenever that member is added to an active zone.

## Protocol intermix

An intermix of FICON and Fibre Channel Protocol (FCP) is supported by FICON Directors at the port level (each port can run either FICON or FCP). This means that the SAN infrastructure can be shared by both protocols to create a commonly managed infrastructure.

## Management capabilities

FICON Directors can be managed using several communication methods, such as:

- ▶ In-band management using a FICON Channel (using the CUP)
- ▶ Out-of-band management using the following:
  - IP-based client/server management software
  - Simple network management protocol (SNMP)
  - Trivial file transfer protocol (TFTP) to load firmware
  - Command-line Interface (CLI) using Telnet
  - GUI provided by the Director vendor or other vendors
- ▶ CLI using the serial interface (RS232)
- ▶ Call home using a modem connection (for notification purposes)

Figure 3-8 on page 56 shows the various interfaces that can be used to manage the FICON Director.

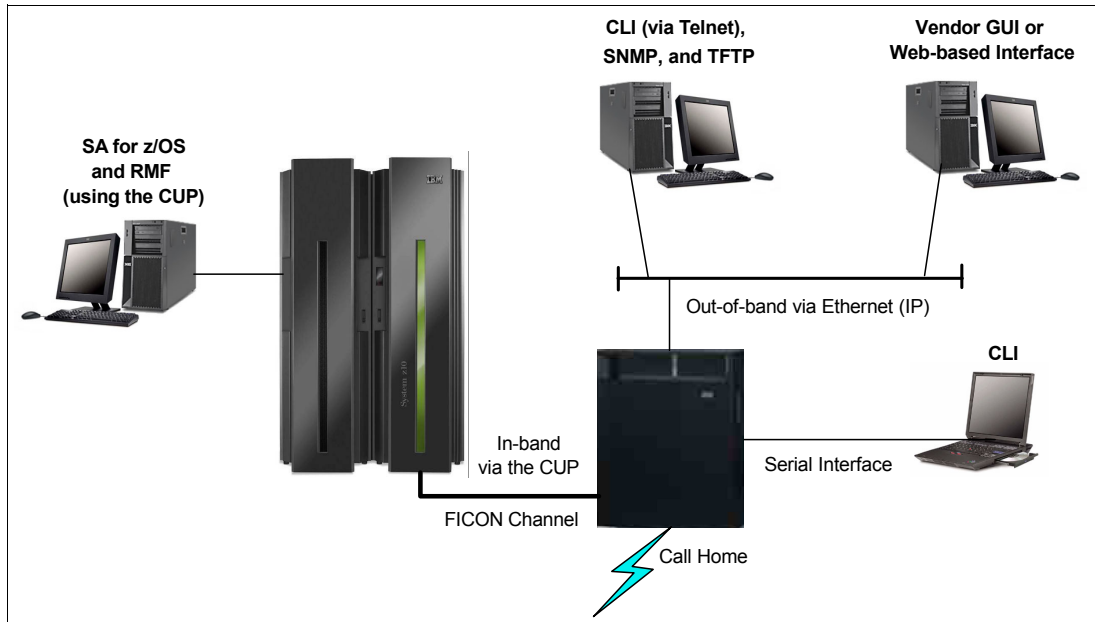


Figure 3-8 FICON management interfaces

## 3.2 Qualified FICON Directors

The purpose of IBM qualification testing for FICON Directors is to ensure that any issues related to interoperability with System z servers (such as functionality, availability, data integrity, and performance) are identified and resolved before the FICON Directors are qualified for System z use.

Qualification testing is carried out in the IBM Vendor Solutions Connectivity (VSC) Lab in Poughkeepsie, New York, USA. The test facilities provide the capability to configure a fully functional FICON environment to drive the various components and protocols of the FICON Directors and simulate failure and recovery actions.

The recovery tests check for the correct detection of a fault or error condition by the attached FICON Directors and ensure that the recovery adheres to System z architecture rules. For a test to be classified successful, all errors must be detected and properly handled by the attached FICON Directors.

In addition, IBM proprietary software and microcode utility test suites stress the FICON Directors to ensure they can perform at levels beyond typical customer usage.

In this section, we focus on seven qualified FICON Directors that support a link data rate of 16 and 8 Gbps: Four from the IBM System Storage SAN b-type family and three from the Cisco MDS 9500 Series.

The SAN b-type family consists of the IBM SAN768B-2, IBM SAN384B-2, IBM SAN768B, IBM SAN384B. The Cisco MDS 9500 Series encompasses the Cisco MDS 9513, Cisco MDS 9509, and the Cisco MDS 9506.

Table 3-1 on page 57 lists the IBM and Cisco FICON Directors and some of their key capabilities.



Table 3-1 IBM and Cisco FICON Directors

| Product                      | IBM Model | Supported protocols | MIX Protocol | Number of ports per Chassis   | Port speed (Gbps)                | ISL speed (Gbps)                 |
|------------------------------|-----------|---------------------|--------------|---|----------------------------------|----------------------------------|
| IBM System Storage SAN768B-2 | 2499-816  | FICON, FCP          | Yes          | 64-512@8 G <sup>a</sup><br>32-384@16 G<br>0-22 FCIP                 | 2/4/8/10/16 FC<br>1/10 GigE FCIP | 2/4/8/10/16 FC<br>1/10 GigE FCIP |
| IBM System Storage SAN384B-2 | 2499-416  | FICON, FCP          | Yes          | 64-256@8 G <sup>a</sup><br>32-192@16 G<br>0-22 FCIP                 | 2/4/8/10/16 FC<br>1/10 GigE FCIP | 2/4/8/10/16 FC<br>1/10 GigE FCIP |
| IBM System Storage SAN768B   | 2499-384  | FICON, FCP          | Yes          | 64-512@8 G <sup>a</sup><br>32-384@8 G<br>0-22 FCIP<br>0-48 CEE/FCoE | 1, 2, 4, 8                       | 2, 4, 8, 10                      |
| IBM System Storage SAN384B   | 2499-192  | FICON, FCP          | Yes          | 32 to 192   | 1, 2, 4, 8                       | 2, 4, 8, 10                      |
| Cisco MDS 9513               | 2054-E11  | FICON, FCP          | Yes          | 12 to 528   | 1, 2, 4, 8                       | 2, 4, 8, 10                      |
| Cisco MDS 9509               | 2054-E07  | FICON, FCP          | Yes          | 12 to 336   | 1, 2, 4, 8                       | 2, 4, 8, 10                      |
| Cisco MDS 9506               | 2054-E04  | FICON, FCP          | Yes          | 12 to 192   | 1, 2, 4, 8                       | 2, 4, 8, 10                      |

a. Not Qualified for FICON environment

To find other System z qualified FICON Directors and Switches, go to:

<http://www-03.ibm.com/systems/networking/switches/san/index.html>

### 3.2.1 IBM System Storage SAN b-type family components

The SAN768B-2 and the SAN384B-2 are built on a common 16 Gbps platform. The SAN768B-2 can switch up to 512 FC ports or up to 384 FICON ports. Switching greater than 256 ports requires the virtual fabric feature (VF) to be configured. The SAN384B-2 has half as many user ports and half as much switching bandwidth as a SAN768B-2. The SAN384B-2 supports up to 256 FC ports or up to 192 FICON ports. The SAN768B and SAN384B are the 8 Gps equivalent directors to the SAN768B-2 and SAN384B-2.

The SAN768B-2 and SAN768B utilize a vertical chassis and measures 14U tall. The chassis has 12 slots. Four slots are utilized for two control processor (CP) blades and two core switching blades. The eight remaining slots can be used for user selectable blade types. The SAN384B-2 and SAN384B utilize a horizontal chassis and measures 8U plus a 1U tall exhaust shelf. The chassis has eight slots. Four slots are utilized for two control processor (CP) blades and two core switching blades. The four remaining slots can be used for user selectable blade types.

The SAN768B-2 and SAN384B-2 slots for user selectable blades can be populated with:

- ▶ 32-port or 48-port 16 Gbps blades. Eight ports can be configured to support 10G FC (typically used for ISLs) on this blade.
- ▶ IP Extension with twelve 8 Gbps-capable FC ports and 12 1G IP ports or 2 10G IP connections.

- ▶ Encryption and CEE/FCoE blades are not supported on chassis used for FICON.
- ▶ 64-port 8 Gbps blades (for FCP connections only—no FICON support)
- ▶ Six-port 10 Gbps blade for dark fiber connections, such as ISLs

The SAN768B and SAN384B slots for user selectable blades can be populated with:

- ▶ 16-port, 32-port or 48-port 8 Gbps blades
- ▶ IP Extension with twelve 8 Gbps-capable FC ports and 12 1G IP ports or 2 10G IP connections
- ▶ Six-port 10 Gbps blade for dark fiber connections (typically used for ISLs)
- ▶ 64-port @8 Gbps blades (for FCP connections only—no FICON support)
- ▶ Encryption and CEE/FCoE blades are not supported on chassis used for FICON.

The SAN768B-2 is based on 256 ports running at 16 Gbps (link data rate) and no oversubscription if all 256 ports are running at 16 Gbps and that traffic travels over the backplane. However, any ingress/egress traffic moving through the ASIC does not cross the backplane (local switching), thereby eliminating oversubscription on configurations larger than 256 ports at 16Gbps.

The SAN384B-2 is based on 128 ports running at 16 Gbps (link data rate) and no oversubscription if all 128 ports are running at 16 Gbps, and that traffic travels over the backplane. However, any ingress/egress traffic moving through the same ASIC does not cross the backplane (local switching), thereby eliminating oversubscription on configurations larger than 128 ports at 16Gbps.

The SAN768B is based on 256 ports running at 8 Gbps (link data rate) and no oversubscription if all 256 ports are running at 8 Gbps and that traffic travels over the backplane. However, any ingress/egress traffic moving through the same ASIC does not cross the backplane (local switching), thereby eliminating over subscription on configurations larger than 256 ports at 8 Gbps.

The SAN384B is based on 128 ports running at 8 Gbps (link data rate) and no oversubscription if all 128 ports are running at 8 Gbps and that traffic travels over the backplane. However, any ingress/egress traffic moving through ASIC does not cross the backplane (local switching), thereby eliminating oversubscription on configurations larger than 128 ports at 8 Gbps.

Table 3-2 through Table 3-8 on page 63 lists the component options that are available with the SAN768B and SAN384B FICON Directors.

Table 3-2 lists the available SAN768B-2, SAN384B-2, SAN768B, or SAN384B blade/hardware options and the associated feature codes.

*Table 3-2 SAN768B-2, SAN384B-2, SAN768B, or SAN384B blade/hardware options*

| Blade/Hardware options                            | Feature code | SAN768B-2 or SAN384B-2 | SAN768B or SAN384B |
|---|--------------|------------------------|--------------------|
| FC 16 Gbps, 48-port switch blade <sup>a b c</sup> | 3648         | X                      |                    |
| FC 16 Gbps, 32-port switch blade <sup>a b c</sup> | 3632         | X                      |                    |
| FC 8 Gbps, 64-port switch blade <sup>d e j</sup>  | 3864         | X                      | X                  |

| Blade/Hardware options   | Feature code | SAN768B-2 or SAN384B-2 | SAN768B or SAN384B |
|--|--------------|------------------------|--------------------|
| FC 8 Gbps Extension blade -<br>12x 8 Gbps FC ports,<br>12 x 1 GigE FCIP ports.<br>Optional 2x 10 GigE ports with FC7892 <sup>f</sup> | 3890         | X                      | X                  |
| FC 8 Gbps, 16-port switch blade <sup>g</sup>   | 3816         |                        | X                  |
| FC 8 Gbps, 32-port switch blade <sup>g</sup>   | 3832         |                        | X                  |
| FC 8 Gbps, 48-port switch blade <sup>g</sup>   | 3848         |                        | X                  |
| FC 10 Gbps 6-ports   | 3870         |                        | X                  |
| FC 4 Gbps Routing blade -<br>16x 4 Gbps FC ports,<br>2x GigE FCIP ports.<br>FC7887 required for optional FCIP <sup>k</sup>           | 3850         |                        | X                  |
| 2 Additional Power Supplies.<br>Required in some setups if using 120 V AC <sup>h</sup>   | 7880         | X                      | X                  |
| <b>These Blades are NOT supported on any chassis with FICON chassis even if in their own logical switch</b>                          |              |                        |                    |
| FCoE 10GbE blade - 48 x 10GbE CEE/FCoE ports. <sup>i f j</sup>   | 3880         | X                      | X                  |
| Encryption Blade;<br>16 x 8 Gbps FC ports<br>96 Gbps of Encryption Processing <sup>j k l</sup>                                       | 3895         | X                      | X                  |

- a. 16 Gbps port blades support E, F, M, EX, D-ports and can operate at 2, 4, 8, 10, or 16 Gbps speeds. The first eight ports of the 16 Gbps port blades can be enabled for 10 Gbps FC when FC7892 is activated. 10 Gbps SFP+ are required
- b. 10 G FC integration on the same port for CWDM/DWDM metro connectivity on the same switch.
- c. Diagnostic Port (D-Port) feature for physical media diagnostic, troubleshooting, and verification services.
- d. Contact an IBM Representative for cabling solutions options.
- e. 64-port blade is populated and shipped with 64 8 Gbps shortwave mSFP fiber optic transceivers.
- f. Four Max per SAN384B-2, SAN192B-2, SAN384B, or SAN192B.
- g. 8 Gbps supports E, F, M, EX-ports and can operate at 1, 2, 4, or 8 Gbps speeds
- h. An additional two power supplies are required when using application blades, even in a 240 VAC environment.
- i. Cannot add any other intelligent blades.
- j. Not Qualified for FICON environment
- k. Two max per SAN384B-2, SAN192B-2, SAN384B, or SAN192B
- l. At publishing time for this book, in-flight encryption and compression on ISLs was not qualified for FICON. Check the FOS release notes for updates on the status of supporting this feature.

Table 3-3 lists the available SAN768B-2, SAN384B-2, SAN768B, and SAN384B transceiver options and the associated feature codes.

*Table 3-3 ISAN768B-2, SAN384B-2, SAN768B, or SAN384B transceiver options.*

| Transceiver options          | Feature code |
|------------------------------|--------------|
| SFP Transceiver 1 GbE Copper | 2551         |

| <b>Transceiver options</b>                          | <b>Feature code</b> |
|---|---------------------|
| FC 16 Gbps SW SFP+ (BR-branded)                     | 2601                |
| FC 16 Gbps SW SFP+ 8-pack (BR-branded)              | 2608                |
| FC 10 Gbps SW SFP+ (BR-branded)                     | 2609                |
| FC 10 Gbps LW SFP+ (BR-branded)                     | 2610                |
| FC 16 Gbps 10 km LW SFP+ (BR-branded)               | 2611                |
| FC 16 Gbps 10 km LW SFP+ 8-pack (BR-branded)        | 2618                |
| FC 8 Gbps SW SFP+ (BR-branded)                      | 2802                |
| FC 8 Gbps SW SFP+ 8-pack (BR-branded)               | 2809                |
| FC 8 Gbps 10 km LW SFP (BR-branded)                 | 2821                |
| FC 8 Gbps 10 km LW SFP 8-pack (BR-branded)          | 2828                |
| FC 8 Gbps 25 km ELW SFP (BR-branded)                | 2881                |
| SFP+ Transceiver 10 Gbps SR                         | 2117                |
| SFP+ 10 Gbps SR 8-pack                              | 2118                |
| SFP+ Transceiver 10 Gbps LR                         | 2151                |
| SFP+ 10 Gbps LR 8-pack                              | 2158                |
| 4 Gbps Short Wave SFP Transceiver <sup>a</sup>      | 2401                |
| 4 Gbps WSW 8-Pack SFP Transceiver <sup>a</sup>      | 2408                |
| 4 Gbps 10 km Long Wave SFP Transceiver <sup>a</sup> | 2411                |
| 4 Gbps 10 km LW 8-Pack SFP Transceiver <sup>a</sup> | 2418                |
| 4 Gbps 4 km Long Wave SFP Transceiver <sup>a</sup>  | 2441                |
| 4 Gbps 4 km LW SFP Transceiver 8 Pack <sup>a</sup>  | 2448                |
| 4 Gbps 30 km ELW SFP Transceiver <sup>a</sup>       | 2480                |
| 10 Gbps FC SW XFP Transceiver                       | 2510                |
| 10 Gbps FC LW XFP Transceiver                       | 2520                |
| 1 GbE Copper SFP Transceiver                        | 2550                |
| 8 Gbps Short Wave SFP Transceiver                   | 2801                |
| 8 Gbps SW SFP Transceiver 8 Pack                    | 2808                |

a. Not applicable for 16G Blades

Table 3-4 on page 61 lists the available Software/License options and associated feature codes.

Table 3-4 SAN768B-2, SAN384B-2, SAN768B, or SAN384B Software/License options.

| Software/License options  |      |
|---|------|
| <ul style="list-style-type: none"> <li>▶ One ICL POD enables half of the ICL ports on core blades of a 2499-816. A second POD feature code enables all ICL ports on a 2499-816 platform.</li> <li>▶ One ICL POD feature code enables all of the ICL ports on the core blades of a 2499-416 platform.<sup>a</sup></li> </ul> | 7871 |
| Integrated Routing: Provides Fibre Channel Routing (FCR) service on all 8 Gbps ports. <sup>b c</sup>  | 7889 |
| Integrated Routing: Provides Fibre Channel Routing (FCR) service on all 8 Gbps ports. <sup>b d</sup>  | 7890 |
| 8 Gbps Routing Advanced Extension License: Enables FCIP Trunking & Adaptive Rate Limiting. <sup>e</sup>   | 7891 |
| 10GbE Activation license enables up to 2 x 10 GbE ports on the FC3890 (FX8-24). <sup>f</sup>  | 7892 |
| Encryption 96Gbps Disk Performance Upgrade Activation—upgrade from 48Gbps Encryption Processing—Requires FC3895.  | 7895 |
| FICON w/CUP: Enables in-band management of FICON. <sup>c</sup>  | 7886 |
| FICON w/CUP: Enables in-band management of FICON. <sup>d</sup>  | 7884 |
| Inter-Chassis License. Requires FC7870. Requires FC7870. <sup>g</sup>   | 7885 |
| SAN384B Inter-Chassis License. Requires FC7870. <sup>h</sup>  | 7883 |
| FCIP/FC High-Performance Extension. Optional for FC. Provides FCIP, hardware-based IP compression and encryption, extensive buffering, FastWrite over FC and FCIP and Tape Pipelining. Only applicable when FC3850 is installed.  | 7887 |
| FICON Accelerator for FC3850 (FR4-18i) only: Enables FICON emulation and read/write tape pipelining over FCIP.  | 7888 |
| Half Bandwidth, Inter-Chassis License. Requires FC7870. <sup>i</sup>  | 7882 |
| FICON Accelerator for FC3890 (FX8-24) only: Enables FICON emulation and read/write tape pipelining over FCIP. <sup>j k</sup>  | 7893 |
| Server Application Optimization: Enhances overall application performance and virtual machine scalability. This is offered for DCX users if they did not receive it in early orders. Adaptive Networking is required and is included in the Base S/W Licenses.  | 7910 |
| Two Additional Power Supplies. Required in some setups if using 120 V AC.   | 7880 |
| Inter-Chassis Cable Kit. Includes (2) ICL cables. <sup>l</sup>  | 7870 |
| Extended Fabrics: Extends support for FC connectivity over 10 km.   | 7803 |

a. Each feature code ordered ships sixteen 64 Gbps QSFP transceivers.

b. Not Qualified for FICON environment.

c. Applicable on 2499-816 or 2499-384.

d. Applicable on 2499-416 or 2499-192.

e. Only applicable to the FC3890 (FX8-24).

f. FC7892 also allows the first eight ports of the 16 Gbps port blade to be enabled at 10 Gbps FC. A 10 Gbps FC transceiver is required.

- g. This is a full ICL license. Order one of these for each 2499-384 when using ICLs to cascade two 2499-384 directors.
- h. Order one of these for each 2499-192 whenever using ICLs to cascade a 2499-192 to either a 2499-192 or a 2499-384 using ICLs.
- i. This is a half ICL license. Order one of these for the 2499-384 when using ICLs to cascade to a 2499-192. FC 7883 must be ordered for the 2499-192.
- j. On 2499-384 requires FC7886 & FC7891.
- k. On 2499-192 requires FC7886 & FC7887.
- l. This is a cable kit consisting of two cables. Order one of these for each director, either 2499-384 or 2499-192

Table 3-5 lists the available fiber optic options and associated feature codes.

*Table 3-5 SAN768B-2, SAN384B-2, SAN768B, or SAN384B fiber optic cable options*

| <b>Fiber optic cable options <sup>a</sup></b>      | <b>Feature code</b> |
|--|---------------------|
| Fiber Cable LC/LC 5 meter 50 µm multimode          | 5305                |
| Fiber Cable LC/LC 25 meter 50 µm multimode         | 5325                |
| LC/LC Fiber Cable 10m 50u multimode (OM3)          | 5810                |
| LC/LC Fiber Cable 10m 50u multimode (OM3) 4-pack   | 5818                |
| Fiber Cable LC/LC 5 meter 50 µm multimode 4-Pack   | 5405                |
| Fiber Cable LC/LC 25 meter 50 µm multimode 4-Pack  | 5425                |
| Fiber Cable LC/LC 31 meter 9 µm single mode 4-Pack | 5444                |
| Fiber Cable LC/LC 31 meter 9 µm single mode        | 5731                |

a. Maximum of 96 four-packs or 255 singles

Table 3-6 lists the available rack mounting options and associated feature codes.

*Table 3-6 SAN768B-2, SAN384B-2, SAN768B, or SAN384B rack mounting options*

| <b>Rack mounting options</b>   | <b>Feature code</b> |
|--|---------------------|
| IBM TotalStorage SAN Cabinet C36   | 9250                |
| Provides plant installation of first 2499-816 or 2499-768 in 2109-C36        | 9281 <sup>a</sup>   |
| Provides plant installation of a second 2499-816 or 2499-768 in 2109-C36     | 9282 <sup>a</sup>   |
| Provides plant hardware for field installation in existing customer 2109-C36 | 9283 <sup>a</sup>   |
| Provides plant installation of first 2499-416 or 2499-384 in 2109-C36        | 9286 <sup>b</sup>   |
| Provides plant installation of a second 2499-416 or 2499-384 in 2109-C36     | 9287 <sup>b</sup>   |
| Provides plant hardware for field installation in existing customer 2109-C36 | 9289 <sup>b</sup>   |
| Standalone Installation. Requires FC9800, 9802, 9831, or 9835                | 9284                |

a. Applicable on 2499-816 or 2499-384.

b. Applicable on 2499-416 or 2499-192.

Table 3-7 on page 63 lists the available power cord options and associated feature codes.

Table 3-7 SAN768B-2, SAN384B-2, SAN768B or SAN384B power cord options

| Power cord options  | Feature code |
|---|--------------|
| Linecord for US, Canada, LatAm, Japan. Required for FC9284. | 9800         |
| Linecord for EMEA. Required for FC9284.                     | 9802         |
| Linecord for Australia, NZ. Required for FC9284.            | 9831         |
| Linecord for North and South Korea. Required for FC9284.    | 9835         |

Table 3-8 lists the available firmware upgrade options and associated feature codes.

Table 3-8 SAN768B-2 or SAN384B-2 firmware upgrade options

| Firmware upgrade options                 | Feature code |
|--|--------------|
| 2499-816 FW Upgrade 1 year <sup>a</sup>  | 7710         |
| 2499-816 FW Upgrade 2 years <sup>a</sup> | 7711         |
| 2499-416 FW Upgrade 1 year <sup>b</sup>  | 7714         |
| 2499-416 FW Upgrade 2 years <sup>b</sup> | 7715         |

a. Applicable on 2499-816

b. Applicable on 2499-416

### 3.2.2 IBM System Storage SAN b-type family functions

This section discusses platform-specific functions that are available on System Storage SAN b-type family Directors for use in FICON environments. The platform-specific functions complement the functions that are described in 3.1.4, “Basic functions of a FICON Director” on page 50.

#### IBM Network Advisor

IBM System Storage Network Advisor is a comprehensive network management application that enables end-to-end management of data center fabrics. IBM Network Advisor features easy-to-use administrative tools that streamline and automate repetitive tasks.

IBM Network Advisor is designed for unified management of data center fabrics, from ports on storage devices all the way to the FICON channel of the System z server. It can be used to configure System Storage SAN b-type family Directors.

IBM Network Advisor provides multiprotocol networking support for:

- ▶ Fibre Channel (FC)
- ▶ Fibre Connection (FICON)
- ▶ Fibre Channel over IP (FCIP)
- ▶ Fibre Channel Routing (FCR)
- ▶ Internet SCSI (iSCSI)

#### Inter-Chassis Link

An Inter-Chassis Link (ICL) is a licensed feature that interconnects two System Storage SAN b-type family Directors. ICL ports in the core blades are used to interconnect the two Directors, potentially increasing the number of usable ports. The ICL ports are internally managed as E\_Ports.

When two System Storage SAN b-type family Directors are interconnected by ICLs, each chassis still requires a unique Domain ID and is managed as a separate Director.

## **Adaptive Networking**

Adaptive Networking is a suite of tools and capabilities that enable you to ensure optimized behavior in the SAN. Even under the worst congestion conditions, the Adaptive Networking features can maximize fabric behavior and provide necessary bandwidth for high priority, mission-critical applications and connections.

The following features are part of the Adaptive Networking suite

- ▶ Bottleneck detection
- ▶ Traffic Isolation Zoning (TI)
- ▶ QoS SID/DID Traffic Prioritization
- ▶ Top Talkers
- ▶ Ingress Rate Limiting

### ***Bottleneck detection***

The bottleneck detection feature identifies devices attached to the fabric that are slowing down traffic.

### ***Traffic Isolation Zoning***

Traffic Isolation (TI) zoning allows you to direct traffic to certain paths in cascaded environments. For this reason, ISLs (E\_Ports) must be included in the TI zone.

TI zones are a logical AND with classic zoning (port or WWNN zoning). TI zoning does *not* replace port or WWNN zoning.

TI zoning has two modes of operation:

- ▶ Failover enabled

With failover enabled, traffic takes the path specified by the TI zone regardless of other paths that are available. If all paths in the TI zone fail, traffic is routed to another TI zone.
- ▶ Failover disabled

With failover disabled, true traffic isolation is achieved. Traffic only takes the path specified by the TI zone. If all paths in the TI zone fail, traffic stops and is not rerouted to another path, even if other paths are available outside the TI zone.

### ***QoS SID/DID Traffic Prioritization***

QoS SID/DID Traffic Prioritization allows you to categorize the traffic flow between a host and target that has a high or low priority. QoS SID/DID Traffic Prioritization requires a Networking license for 8 Gbps platforms but does not require a license for 4 Gbps platforms.

### ***Top Talkers in Advanced Performance Monitoring***

The Top Talkers feature is part of the licensed Advanced Performance Monitoring feature. It measures and ranks top bandwidth-consuming traffic in real time based on SID/DID pairs. Top Talkers can identify the ports that consume the most bandwidth and that information can be used to assist you with configuring the SID/DID pairs with certain QoS attributes so they get proper priority.

### ***Ingress rate limiting***

Ingress rate limiting is a licensed feature that requires the Adaptive Networking license. It can restrict the speed of traffic from a particular device to the Director port. Ingress rate limiting



can also be used to reduce existing congestion in the network or proactively avoid congestion.

### **Frame-level trunking**

Frame-level trunking automatically distributes data flows over multiple physical Inter-Switch Link (ISL) connections and logically combines them into a trunk to provide full bandwidth utilization while reducing congestion. Frame-level trunking can:

- ▶ Optimize link usage by evenly distributing traffic across all ISLs at the frame level
- ▶ Maintain in-order delivery to ensure data reliability
- ▶ Help ensure reliability and availability even if a link in the trunk fails
- ▶ Optimize fabric-wide performance and load balancing with Dynamic Path Selection (DPS)
- ▶ Simplify management by reducing the number of ISLs required
- ▶ Provide a high-performance solution for network- and data-intensive applications

Up to eight ISLs can be combined into a trunk, providing up to 128 Gbps data transfers (with 16 Gbps ISLs).

### **Virtual ISL connection**

When a base switch is connected to another base switch, a Virtual ISL (XISL) connection is created. When logical switches with the same fabric ID (FID) are configured to use the XISL (such as the logical switches #1 and #2), the base switches automatically create a Logical ISL (LISL) within the XISL. The LISL isolates traffic from multiple fabrics because each LISL is dedicated to traffic for a single fabric.

### **Lossless Dynamic Load Sharing**

Lossless Dynamic Load Sharing (DLS) allows you to rebalance trunk port paths without causing input/output (I/O) failures in cases where the end devices require in-order delivery (IOD) of frames. You can use DLS on trunks connecting switches to perform the following functions:

- ▶ Eliminate dropped frames and I/O failures by rebalancing the paths going over the ISLs whenever there is a fabric event that might result in sub-optimal utilization of the ISLs
- ▶ Eliminate the delay caused by establishing a new path when a topology change occurs
- ▶ Pause ingress traffic (by not returning credits)
- ▶ Change the existing path to a more optimal path
- ▶ Wait for sufficient time for frames already received to be transmitted
- ▶ Resume traffic after a failure

### **Virtual fabrics or logically partition a physical SAN**

The Virtual Fabrics feature enables a physical storage area network (SAN) switch to be partitioned into multiple logical switches, each of which is assigned to a separate logical fabric. A logical switch and logical fabric are independent entities that are managed just like a physical switch or fabric. The Virtual Fabrics feature provides data, control, and management isolation for logical switches and logical fabrics using a standards-based implementation of the ANSI standard for virtual fabrics.

Vitalization supports the following use cases:

- ▶ Isolating traffic types
- ▶ Pooling resources for independent fabrics
- ▶ Isolating different application tiers and device management
- ▶ Facilitating backbone edge routing
- ▶ Increasing FICON consolidation and scalability

- ▶ Connect logical and “unaware” switches to form logical fabrics
- ▶ Isolate and manage by application, business group, customer, or traffic type (FC, FICON)

Table 3-9 lists the logical switch types supported by the Virtual Fabrics feature.

Table 3-9 Logical switch types

| Switch type    | Description  |
|----------------|--|
| Base switch    | Not supported for FICON.<br>The base switch can only contain ISLs that are shared by other logical switches. Although not supported for FICON, the base switch can be used for other logical switches.   |
| Default switch | The default switch is FID=128, and it is where all ports are placed when virtual fabrics are enabled. Virtual fabrics are enabled by default for all switches shipped from the factory with FOS 6.2.0 and above.<br>Only port-based numbering is permitted in the default switch. Zero-based port numbering is not supported. As a result, the default switch is usually not used for FICON. |
| Logical switch | The logical switch is the most common switch type used for FICON. A logical switch can have any user-defined FID between 1-127. Logical switches support both port-based and zero-based port numbering.  |

### Port numbering

A port number is a number assigned to an external port to give it a unique identifier to the FICON Director. Ports are identified by both the slot number in which the blade is located in the chassis and the port number on the blade (slot number/port number). Table 3-10 shows the port numbering scheme for each blade type.

Table 3-10 Port numbering scheme

| Blade type | Blade port numbering  |
|------------|---|
| 16-port    | Ports are numbered from 0 through 15 from bottom to top.  |
| 32-port    | Ports are numbered from 0 through 15 from bottom to top on the left set of ports, and 16 through 31 from bottom to top on the right set of ports. |
| 48-port    | Ports are numbered from 0 through 23 from bottom to top on the left set of ports, and 24 through 47 from bottom to top on the right set of ports. |
| 6-port     | Ports are numbered from 0 through 5 from bottom to top.   |

Any port can be assigned to any logical or default switch. When zero-based addressing is used, ports can have any number the user defines. Some forward planning, however, can simplify the task of numbering the ports and minimize service impacts.

### Migration

The combination of logical switches and zero-based addressing make for easy migration from IBM System Storage SAN b-type and m-type FICON Directors to a SAN768B or SAN384B platform. For each FICON Director to be migrated, a logical switch with its Domain ID and port numbers identical to those in the old FICON Director can be created. Therefore, you can perform a migration without making any changes to HCD, HCM, or IOCP.

Table 3-11 on page 67 lists some of the functions and their associated hardware, software, and firmware prerequisites that are discussed in this section for the SAN768B and SAN384B.

Table 3-11 Hardware, software, and firmware prerequisites

| Function   | Hardware prerequisites  | Software and firmware prerequisites  |
|--|-------------------------|--|
| IBM Network Advisor                                    | All supported Directors | Refer to:<br><a href="http://www-01.ibm.com/common/ssi/rep_ca/3/897/ENUS211-243/ENUS211-243.PDF">http://www-01.ibm.com/common/ssi/rep_ca/3/897/ENUS211-243/ENUS211-243.PDF</a> |
| Inter-Chassis Link (ICL)                               | SAN768B or SAN384B      | FOS 6.1 and above  |
| Traffic Isolation (TI)                                 | SAN768B or SAN384B      | FOS 6.0.0c and above   |
| Virtual Fabric (logical partition)                     | SAN768B or SAN384B      | FOS 6.2 and above  |
| Quality of Service (QoS)                               | SAN768B or SAN384B      | FOS 6.0.0c and above   |
| Adaptive Network TopTalker, QoS, Ingress Rate Limiting | SAN768B or SAN384B      | FOS 6.0.0c and above   |

For more details regarding the features and functions provided by the System Storage SAN b-type family, refer to the following:

- ▶ *IBM System Storage SAN768B-2 Installation, Service, and User Guide, GA32-0893-00*
- ▶ *IBM System Storage SAN384B-2 Installation, Service, and User Guide, GA32-0894-00*
- ▶ *IBM System Storage SAN48B-5 Installation, Service, and User Guide, GA32-0895-00*
- ▶ *IBM System Storage SAN768B Director Installation, Service and User's Guide, GA32-0574*
- ▶ *IBM System Storage SAN384B Director Installation, Service and User's Guide, GC52-1333*
- ▶ *Brocade Fabric OS Administrator's Guide, 53-1002148-03*
- ▶ IBM System Storage SAN768B-2 and SAN384B-2 web page:  
<http://www-03.ibm.com/systems/networking/switches/san/b-type/san768b-2/index.html>
- ▶ IBM System Storage SAN768B and SAN384B web page:  
<http://www-03.ibm.com/systems/networking/switches/san/b-type/san768b/index.html>
- ▶ IBM System Storage SAN48B-5 web page:  
<http://www-03.ibm.com/systems/networking/switches/san/b-type/san48b-5/index.html>

### 3.2.3 Cisco MDS 9500 Series

The Cisco MDS 9500 Series is a modular multi-layer platform with common switching and service modules. The Cisco MDS 9513 has 11 slots for various modules and can switch up to 528 ports simultaneously. The Cisco MDS 9509 has seven slots for various modules and can switch up to 336 ports simultaneously. The Cisco MDS 9506 has four slots for various modules and can simultaneously switch up to 192 ports.

The Cisco MDS 9500 Series utilizes a horizontal chassis. The Cisco MDS 9513 and MDS 9509 with their dual Supervisor-2 modules measure 14U tall. The Cisco MDS 9506 has dual Supervisor modules and measures 7U tall.

The module slots in all three platforms can be filled with:

- ▶ IP services, storage services, and multi-service modules
- ▶ 12-port, 24-port, and 48-port 4 Gbps switching modules
- ▶ 24-port, 48-port, and 4/44-port 8 Gbps switching modules
- ▶ 4-port 10 Gbps switching module

When all 48 ports on a 48-port switching module are used simultaneously at 8 Gbps, the oversubscription factor can be as much as 4 : 1. On a 24-port switching, when all 24-ports module are used simultaneously at 8 Gbps, the oversubscription factor can be as much as 2 : 1.

Table 3-12 through Table 3-16 on page 71 list the component options available with the Cisco MDS 9500 Series FICON Directors.

Table 3-12 lists the available line card options and associated feature codes.

*Table 3-12 Cisco MDS 9500 Series line card options*

| <b>Line card options</b>                            | <b>Feature code</b> |
|---|---------------------|
| 8-Port IP Services Line Card                        | 2208                |
| Multiprotocol Services Module                       | 2214                |
| Storage Service Module                              | 2400                |
| 4 port 10 Gbps FC Switch Module                     | 2404                |
| 12 port 4 Gbps FC Switch Module                     | 2412                |
| 24 port 4 Gbps FC Switch Module                     | 2424                |
| 48 port 4 Gbps FC Switch Module                     | 2448                |
| 18/4 Multiservice Module <sup>a</sup>               | 2450                |
| 16 Port SSN 16 Module <sup>b</sup>                  | 2816                |
| 32 Port Adv 8G FC Module <sup>c</sup>               | 2818                |
| 48 Port Adv 8G FC Module <sup>c</sup>               | 2820                |
| 24 8 Gbps FC Switching Module <sup>d</sup>          | 2824                |
| 4/44 port 8 Gbps Host Switching Module <sup>d</sup> | 2844                |
| 48 port 8 Gbps FC Switching Module <sup>d</sup>     | 2848                |
| 8 Port 10Gb FCoE Blade <sup>c</sup> SFP             | 2950                |
| Sup 2A Modules 2 Field <sup>c</sup>                 | 3002                |
| 9513 Fabric2 Module - Field                         | 3100                |
| 12-port 4 Gb FC Switch Mod                          | 8S0857              |
| 24-port 4 Gb FC Switch Mod                          | 8S0858              |
| 48 Port 4 GB FC Module                              | 8S0859              |
| 9513 Fabric3 Switching Modules                      | 3101                |

- a. requires OS 3.2.2 above
- b. requires OS 4.2.1 or above
- c. requires OS 5.2 or above
- d. requires OS 4.1.1 or above

Table 3-13 lists the available transceiver CU SFP options and associated feature codes.

Table 3-13 Cisco MDS 9500 Series transceiver and CU SFPs

| Transceiver options   | Feature code |
|---|--------------|
| FC 10Gb FC SW SFP   | 5020         |
| FC 10Gb FC LW SFP   | 5022         |
| FC 10 Gbps SRX2 Transceiver   | 5030         |
| FC Ethernet 10 Gbps SRX2 Transceiver                                      | 5032         |
| FC 10 Gbps 10 km LWX2 SC Transceiver                                      | 5040         |
| FC 10Gbps 40 km ERX2 Transceiver  | 5050         |
| FC 4 Gbps SW SFP Transceiver - 4 Pack                                     | 5434         |
| FC 4 Gbps 4 km LW SFP Transceiver 4 Pack                                  | 5444         |
| FC 4 Gbps 10 km LW SFP Transceiver - 4 Pack                               | 5454         |
| FC 8 Gbps SW SFP+ Transceiver (requires OS 4.1.1 or later)                | 5830         |
| FC 8 Gbps SW SFP+ Transceiver - 4 Pack (requires OS 4.1.1 or later)       | 5834         |
| FC 8 Gbps 10 km LW SFP+ Transceiver (requires OS 4.1.1 or later)          | 5850         |
| FC 8 Gbps 10 km LW SFP+ Transceiver - 4 Pack (requires OS 4.1.1 or later) | 5854         |
| 8 Gbps FC LC 40 km SFP  | 5855         |
| Tri-Rate SW SFP Transceiver   | 5210         |
| Tri-Rate LW SFP Transceiver   | 5220         |
| Gigabit Ethernet Copper SFP (feature code 2450 is required)               | 5250         |
| TRI-RATE SW SFP   | 5210         |
| Tri-Rate LW SFP   | 5220         |
| FC 2 Gb SW SFP  | 5230         |
| FC 2 Gb LW SFP  | 5240         |
| Gigabit Ethernet Copper SFP   | 5250         |
| 10GBASE CU SFP 1m Cable   | 5410         |
| 10GBASE CU SFP 3m Cable   | 5420         |
| 10GBASE CU SFP 5m Cable   | 5430         |
| 10GBASE CU SFP 1m Cable   | 5510         |
| 10GBASE CU SFP 3m Cable   | 5520         |
| 10GBASE CU SFP 5m Cable   | 5530         |

Table 3-14 on page 70 lists the available hardware and software packages and associated feature codes.

Table 3-14 Cisco MDS 9500 Series hardware and software packages

| Hardware and software packages         | Feature code |
|--|--------------|
| FICON package                          | 1000         |
| 9200 FCIP Act/8-Port IP Srv LC         | 2209         |
| 9500 FCIP Act/8-Port IP Srv LC         | 2210         |
| 9200 FCIP Activation for MSM           | 2215         |
| 9500 FCIP Activation for MSM           | 2216         |
| MDS 9500 Enterprise Package            | 7021         |
| MDS 9500 Fabric Manager Server Package | 7026         |
| MDS 9200 Enterprise Pkg                | 7011         |
| MDS 9200 Fab Mgr Serv Pkg              | 7016         |
| MDS 9200 Mainframe Server Pkg          | 7031         |
| MDS 9500 SSE Package                   | 7040         |
| MDS 9200 Storage Serv Enabler          | 7041         |
| MDS 9200 SME on Chassis                | 7044         |
| MDS 9200 SME on 18 4                   | 7045         |
| MDS 9200 SME on SSN 16                 | 7046         |
| MDS 9500 SME on 18 4                   | 7047         |
| MDS 9500 SME on SSN 16                 | 7048         |
| MDS 9500 E delivery DCNM Adv           | 7049         |
| MDS 9200 E delivery DCNM               | 7050         |
| MDS 9200 FCIP Act for 18/4             | 7052         |
| MDS 9500 FCIP Act for 18/4             | 7055         |
| 9200 I/O Acc for 18/4                  | 7062         |
| 9200 I/O Acc. Base Machine             | 7072         |
| 9500 I/O Acc for 18/4                  | 7075         |
| 9200 FCIP Act for SSN16                | 7082         |
| 9500 FCIP Act for SSN16                | 7085         |
| 9200 I/O Acc for SSN16                 | 7092         |
| 9500 I/O Acc for SSN16                 | 7095         |
| Sup2A Plant                            | 9002         |
| 9513 Fabric2 Modules-Plant             | 9100         |
| 9513 Fabric3 Switching Modules         | 9101         |
| MDS 9500 Enterprise Pkg                | 8S0860       |
| MDS 9500 Fabric Mgr Serv Pkg           | 8S0861       |

| Hardware and software packages    | Feature code |
|-----------------------------------|--------------|
| MDS 9500 Mainframe Server Pkg     | 8S0862       |
| MDS 9500 Mainframe Server Package | 7036         |

Table 3-15 lists the available fiber optic cable options and associated feature codes.

*Table 3-15 Cisco MDS 9500 Series fiber optic cable options*

| Fiber optic cable options                    | Feature code |
|--|--------------|
| 5 meter 50 µm LC/LC Fiber Cable (multimode)  | 5605         |
| 25 meter 50 µm LC/LC Fiber Cable (multimode) | 5625         |
| 5 meter 50 µm LC/LC Fiber Cable - 4 Pack     | 5642         |
| 25 meter 50 µm LC/LC Fiber Cable - 4 Pack    | 5643         |
| 1 M LC/LC Fiber Cable                        | 5601         |
| 1 m LC/LC Fiber Cable 4 Pack                 | 5641         |
| Gigabit Ethernet Copper SFP                  | 5250         |
| 10GBASE CU SFP 1m Cable                      | 5410         |
| 10GBASE CU SFP 3m Cable                      | 5420         |
| 10GBASE CU SFP 5m Cable                      | 5430         |
| 10GBASE CU SFP 1m Cable                      | 5510         |
| 10GBASE CU SFP 3m Cable                      | 5520         |
| 10GBASE CU SFP 5m Cable                      | 5530         |

Table 3-16 lists the available rack mounting options and associated feature codes.

*Table 3-16 Cisco MDS 9500 Series rack mounting options*

| Rack mounting options               | Feature code |
|-------------------------------------|--------------|
| Plant Install 9506 in 2109-C36      | 9523         |
| Field Merge 9506 in 2109-C36        | 9524         |
| Plant Install 9509 in 2109-C36      | 9533         |
| Field Merge 9509 in 2109-C36        | 9534         |
| Plant Install 9513 in 2109-C36 rack | 9543         |
| Field Merge 9513 in 2109-C36 rack   | 9544         |

### 3.2.4 Functions of the Cisco MDS 9500 Series

This section discusses platform-specific functions that are available on the Cisco MDS 9500 Series of Directors for use in FICON environments. The platform-specific functions complement the functions that are described in 3.1.4, “Basic functions of a FICON Director” on page 50.

## **Fabric Manager**

The Fabric Manager is a set of network management tools that supports secure Simple Network Management Protocol version 3 (SNMPv3) and earlier versions. It provides a graphical user interface (GUI) that displays real-time views of your network fabric and lets you manage the configuration of Cisco MDS 9500 Series FICON Directors. Detailed traffic analysis is also provided by capturing data with SNMP. The captured data is compiled into various graphs and charts that can be viewed with any web browser.

The Cisco Fabric Manager applications are:

- ▶ Fabric Manager Server
- ▶ Device Manager

### ***Fabric Manager Server***

Fabric Manager Server is a separately licensed feature that extends the capabilities of the basic Fabric Manager. It adds the ability for Fabric Manager to manage multiple fabrics simultaneously (note that in Fibre Channel terms, each Director in switched or cascaded FICON Director topology, is a separate fabric). It also enables accumulation and analysis of historical performance data, which can be a valuable complement to RMF.

### ***Device Manager***

The Device Manager presents two views of a Director:

1. Device View displays a graphic representation of the switch configuration, and provides access to statistics and configuration information for a single switch.
2. Summary View displays a summary of E\_Ports (Inter-Switch Links), F\_Ports (fabric ports), and N\_Ports (attached hosts and storage) on the Director.

### ***Comprehensive security framework***

The Comprehensive security framework supports RADIUS authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), Fibre Channel Security Protocol (FC-SP), VSANs, hardware-enforced zoning, port security, fabric binding and VSAN-based access control.

## **PortChannel (ISL Trunking)**

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

### **Trunk E port (TE port)**

Trunking, also known as VSAN trunking, is a feature that is specific to switches in the Cisco MDS 9500 Series. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

The trunking feature includes:

- ▶ Trunking configurations are applicable to E\_Ports. If trunk mode is enabled on an E\_Port and that port becomes operational, it is referred to as a TE port.
- ▶ The trunk-allowed VSANs configured for TE ports are used by the trunking protocol to determine the allowed-active VSANs in which frames can be received or transmitted.
- ▶ If a trunking enabled E\_Port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E\_Port



## Quality of Service (QoS)

The QoS feature in Cisco SAN-OS allows for data traffic to be classified into distinct levels for service differentiation ( low, medium, or high priority).

Quality of service (QoS) offers the following:

- ▶ Provides a relative bandwidth guarantee to application traffic
- ▶ Controls the latency experienced by application traffic
- ▶ Prioritizes one application over another (for example, prioritizing transactional traffic over bulk traffic) through bandwidth and latency differentiation

You can apply QoS to ensure that Fibre Channel data traffic for your latency-sensitive applications receive higher priority over throughput-intensive applications, such as data warehousing.

## Virtual Storage Area Network

Virtual Storage Area Network (SAN) technology is used for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON intermix support.

The Cisco MDS 9500 Series can have up to 256 unique VSANs, and eight of those can be used for FICON. Each VSAN has its own set of fabric services (fabric server, name server, and so on), FICON Control Unit Port, Domain ID, Fabric Shortest Path First (FSPF) routing, operating mode, IP address, and security profile.

VSANs can span line cards and are dynamic in size. For example, one VSAN with 10 ports can span 10 different line cards. Adding ports to a VSAN is a nondisruptive process. The maximum number of ports for a VSAN is 255 due to FICON addressing rules.

VSANs enable complete traffic isolation and are maintained throughout the FICON fabric. You can also assign specific VSANs to certain cascaded links.

### ***VSAN-Based Access Control***

This feature enables customers to define roles where the scope of the roles is limited to certain VSANs. For example, a network administrator role can be set up to allow configuration of all platform-specific capabilities, and VSAN-administrator roles can be set up to only allow configuration and management of specific VSANs.

VSAN-based access control reduces SAN disruptions by localizing the effects of user errors to the VSANs for which the user has administrative privileges. It adds a layer of security where only administrators can configure switches within specified VSANs.

## Port numbering

A range of 250 port numbers is available for you to assign to all the ports on a Director. You can have more than 250 physical ports assigned and the excess ports do not have port numbers in the default numbering scheme. You can have ports without a port number assigned if they are not in a FICON VSAN, or you can assign duplicate port numbers if they are not used in the same FICON VSAN. By default, port numbers are the same as port addresses, and the port addresses can be swapped.

The following rules apply to FICON port numbers:

- ▶ Supervisor modules do not have port number assignments.
- ▶ Port numbers do not change based on TE ports. Because TE ports appear in multiple VSANs, chassis-wide unique port numbers are reserved for TE ports.

- ▶ Each PortChannel must be explicitly associated with a FICON port number.
- ▶ When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- ▶ Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, then the associated ports will not come up.

## **Migration**

The combination of VSANs and FICON port numbers make for easy migration from previous-generation FICON Directors to a Cisco MDS 9500 infrastructure. For each FICON Director to be migrated, a FICON VSAN with its Domain ID and FICON port numbers identical to those in the old FICON Director can be created. Therefore, you can perform a migration without making any changes to IOCP or HCD.

For more details regarding the features and functions provided by the Cisco MDS 9500 Series, refer to the following web site:

<http://www-03.ibm.com/systems/networking/switches/san/ctype/9500/index.html>



## Part 2

# Planning the FICON environment

This part provides an overview of the items that you must consider when planning the FICON environment.

It provides a structured approach for planning and designing an end-to-end FICON infrastructure solution. It also gives a solid base for configuring the components of the solution during the implementation phase.





## Planning the FICON environment

When you are building a communication infrastructure with the highest quality of service for a FICON environment, proper planning is crucial.

This chapter provides you with a structured approach for planning and designing an end-to-end FICON infrastructure solution. It gives a solid base for configuring the components of the solution during the implementation phase.

## 4.1 Structured approach for planning

Having a structured approach for planning is key to reaching your desired goal when you are designing any end-to-end solution. The approach used in this chapter offers a framework that is based on best practices and industry standards.

The topics we cover follow a logical sequence that aids in the decision-making process for planning a FICON environment. The decision points are listed here and described in greater detail in the following pages:

- ▶ Documentation
- ▶ Requirements
- ▶ Context
- ▶ Topologies and supported distances
- ▶ Convergence
- ▶ Management
- ▶ Virtualization and availability
- ▶ Performance
- ▶ Prerequisites and interoperability
- ▶ Physical connectivityRequirements

If you are not familiar with Fibre Channel and FICON architectures and terminology, before proceeding we strongly recommend that you read Chapter 1, “Introduction to FICON” on page 3; Chapter 2, “System z FICON technical description” on page 21; and Chapter 3, “FICON Director technical description” on page 45.

In addition, to be successful in designing a FICON solution, you must be well versed in the technologies and products that support typical FICON solutions. Although this book provides some guidance, it is not meant to provide in-depth technology and product knowledge. Therefore, we recommend that you also attend the appropriate training and education offered by IBM as described at the following URL:

<http://www-304.ibm.com/jct03001c/services/learning/ites.wss/zz/en?pageType=page&c=a0001742>

### Documentation

Current and accurate documentation (planning, design, installation, and final) is essential for a successful implementation and later operations of the FICON environment. However, because most information is not always readily available during the various phases of the planning process, there are situations where it requires iterations of the process.

To avoid delays during the planning and implementation process, ensure that the various types of documentation needed are always well-defined, complete, and current.

### Requirements

Planning for any end-to-end solution is a process of iterations. It starts with gathering requirements and defining the desired outcome of the solution.

After all relevant data is collected, you must review and evaluate it and outline the objectives of the solution.

If you are migrating from an existing environment, you must also consider the current configuration, including the functions and features that are implemented and how they are being used.

## Context

The output of the requirements-gathering phase must provide a solid base for mapping to one or more of the following scenarios:

- ▶ Migrating from ESCON to FICON
- ▶ Moving to an 8 Gbps environment (FICON Express8/8s)
- ▶ Migrating from a single site to a multi-site environment
- ▶ Implementing a new FICON environment

## Topologies and supported distances

To fulfill the requirements of the chosen solution (or scenario), you must select the most suitable topology. The total end-to-end distance (from a channel to a control unit) is a main factor in determining the correct topology.

The following topologies are supported:

- ▶ Point-to-point
- ▶ Switched
- ▶ Cascaded

For extended distances, you must also evaluate the IBM qualified solutions depending on your requirements and the supported distances of the FICON connections.

Note that IBM owns a strict qualification process for giving assurance of an extended distance solution based on WDM technology.

## Convergence

The I/O communication infrastructure (FICON fabric) provides the ability to attain high availability and continuous operations in a predictable manner, with centralized management and scalability for future changes. This means that the switching platform must be robust enough to deliver such capability and be ready to deploy new enhancements that improve throughput, security, and virtualization. Therefore, the FICON Director is the reference platform for FICON fabrics.

Factors that influence the selection of the FICON Director platform are:

- ▶ Intermix of the FICON and Fibre Channel Protocol (FCP), which enables you to consolidate the I/O infrastructure, maximize the utilization of the bandwidth, and consolidate the fabric management.
- ▶ Security, high integrity, and zoning are integral parts of the entire FICON infrastructure.

Other relevant factors for the FICON fabric, such as power consumption, cooling, and space inside a data center, must also be considered.

## Management

The selection of the systems management platforms and their integration with the operation support environment is important to the continuous operations of the FICON infrastructure.

The possibilities for FICON management include:

- ▶ Command-line Interface
- ▶ Element management
- ▶ Fabric management
- ▶ Storage management initiative specification
- ▶ System z management for FICON Directors

## Virtualization and availability

Because the efficiency and effectiveness of a data center is determined by the levels of virtualization and availability, the complete infrastructure must be considered. All aspects of virtualization and physical connectivity for the FICON channel, Director, and control unit must be evaluated.

## Performance

The selection of the right technology is determined by satisfying traffic patterns, further segmentation, and performance of the end-to-end solution. It is important to understand the fan-in, fan-out, and over subscription ratios for the entire FICON environment. Traffic management gives you the possibility to proactively manage any congestion inside the fabric and minimize latency. There are a number of design evaluation tools that can aid in this process.

## Prerequisites and interoperability

Checking the compliance of all FICON components in the infrastructure with the prerequisites of the planned solution is key to ensuring that your implementation will proceed without delays and avoid interoperability problems. The FICON components include the System z FICON features and control units and the FICON Directors and Wavelength Division Multiplexor (WDMs) platforms, if applicable to the selected topology.

Many failures or inter-operability and performance issues can be put down to microcode levels and OS levels being below required levels. We recommend ensuring that microcode levels are checked to ensure they are at the latest level (including other vendors) and that APARs are checked to ensure they all the relevant APARs are installed for FICON.

- ▶ System z: Apply latest driver level and apply latest MCLs.
- ▶ Director: Apply latest FOS.
- ▶ Storage: Apply latest microcode/firmware.
- ▶ zOS: Apply all FICON related APARS.

Further information can be found at the System Storage Interoperation Center (SSIC)

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

## Physical connectivity

The physical transport media must be selected according to the installed or planned features. Choosing a qualified cabling solution based on industry standards for Data Center Cabling (US TIA-942/Europe EN50173-5) and best practices in accordance with System z qualified IBM Facilities Cabling Services will provide a secure and highly scalable consolidated physical layer for the complete data center.

In the subsequent sections, we describe considerations and recommendations for each of the identified decision points in more detail.

## 4.2 Documentation

Planning, design, installation, and final documentation are requirements for a successful implementation and later operations of the FICON environment. Based on best practices and requirements we recommend the following types of documentation.

Planning documentation:

- ▶ Current environment



- ▶ Requirements
- ▶ Available features and functions
- ▶ Decision points

Design documentation:

- ▶ High level
- ▶ Detailed level

Installation documentation:

- ▶ Detailed implementation plan

Final documentation:

- ▶ Technical
- ▶ Operations
- ▶ Disaster recovery plan
- ▶ Change procedure

The number of documents that you must create depends on the complexity of the solution and the required level of detail.

Creating and maintaining documentation throughout the planning, design, and implementation phases of your FICON infrastructure is an important part of running the environment without unplanned outages.

For the most part, the implementation, operations, and support (troubleshooting) responsibilities are owned by different groups. From a security viewpoint, there is usually a clear boundary between the responsibilities. Changes must be monitored and audited to ensure that due diligence is observed; for example, execution must be strictly separated from auditing. And finally, you must consider the planned and unplanned turnover of professionals at the workplace.

To reduce the risk of disruption to your FICON infrastructure, keep your documentation as current and accurate as possible and store a copy in a secure place, preferably in a different location in case of a disaster.

## 4.3 Requirements

It is important to clearly understand and accurately document all requirements. Such documentation will help you throughout the planning and design process.

After the requirements are collected and documented, each one must be carefully evaluated. For existing environments (FICON or ESCON), it is also important to identify all equipment that is currently installed and how it is being used. This means physical and logical inventories must be carried out.

The goal of the physical inventory is to identify and verify what is installed in your FICON or ESCON environment. Through onsite visual inspections and documentation, all ESCON or FICON channels, Directors and ports, control units, and cabling must be identified.

Although the overall goal of the physical inventory is to identify and verify what you physically installed, the goal of the logical inventory is to understand the functions that are being used and how they are defined. Determine whether the functions must be replaced by a FICON function.

During the planning and design process, the requirements are mapped to features, functions, topology, and technologies, which in turn determines an approximate design of your FICON environment. This does not mean you only have one solution, but rather multiple alternatives. Assess the alternatives using all of the steps in this chapter.

## 4.4 Context

To define the context of your environment, consider the existing and planned components (System z, FICON Director, and control units) that were determined and documented in the requirements gathering phase. Then define your transport layer strategy dependent on the upper-layer requirements for high availability, disaster recovery, and business continuity. This all must be synchronized with the communication layer requirements of your applications.

After all relevant data is reviewed and analyzed, use the output (alternatives) to map to one or more of the following scenarios:

- ▶ Migrating from ESCON to FICON (changing the transport layer protocol)
- ▶ Moving to a high bandwidth environment (implementing FICON Express8/8S)
- ▶ Migrating from a single site to a multi-site cascaded topology (building a highly available multi-site solution)
- ▶ Implementing a new FICON environment (building a completely new FICON environment)

The next sections provide additional information for each scenario.

### 4.4.1 Migrating from ESCON to FICON

This scenario discusses the considerations for migration from an ESCON environment to a FICON environment. We concentrate on DASD migration because it is the most commonly connected device type.

FICON is addressing technical limitations of ESCON in bandwidth, distances, topology, channel and device addressing and performance.

For the right number of FICON channels you need:

- ▶ The I/O operations per second and their characteristics
- ▶ The required bandwidth

ESCON channels are planned with 50% utilization or less and depend on block size. With higher utilization, the traffic is slower. Another limit is the I/O rate, which is typically less than 1000 I/O per second. The ESCON data droop is caused by the single CCW transmit at a time, and then waiting for the notification.

We do not recommend planning over the actual limits because that can cause performance degradation. It is also important to have enough channels to keep the number of concurrent I/Os low.

You also must determine the peak I/O and MB/s values for each subsystem, which reflects current workload and planned growth. We recommend that you use analysis tools over a longer time period to obtain realistic trends.

For native tape, throughput is important. Aggregating multiple ESCON tape channels onto a single FICON channel reduces the System z infrastructure. This can be reached due to high head-to-tape transfer rates, capacity, and accessibility of the cartridge.

Using FICON for virtual tapes also reduces the number of required CUs, back-end tape drives, and channel paths.

When planning the migration from an existing ESCON configuration to a FICON configuration, consider the following rules:

- ▶ Configure at least two channel paths to a logical control unit for high availability. (More channel paths might be required, depending on throughput requirements.)
- ▶ A logical control unit or device cannot be defined to be accessed more than once from the same channel path.
- ▶ A physical control unit that has multiple logical control units (specified by the CUADD parameter in IOCP/HCD) might be accessed more than once from the same FICON (FC) channel path, but the access is to different CUADDs (different logical control units) within the physical control unit.
- ▶ Configure the channel paths according to the quantity of resources available in the FICON channel and control unit.
- ▶ It is not possible to aggregate two or more ESCON channel paths that access the same logical control unit into only one FICON channel path.

Even if there are the same number of paths from the operating system image to the disk subsystem in both ESCON and FICON configurations, the advantages of using FICON are:

- ▶ More concurrent I/Os to the same control unit
- ▶ Concurrent I/Os with other control units
- ▶ Link data rate (200 Mbps for ESCON versus up to 8 Gbps for FICON)
- ▶ Unrepeated distance from channel (3 km for ESCON and 10 km for FICON)

#### 4.4.2 Moving to a high bandwidth environment (FICON Express8/8S)

Moving to 8 Gbps FICON improves the performance, scalability, and efficiency of your System z and storage environment. With the use of a FICON Director you can utilize the physical channels even more when shared with multiple connections to different CUs at the same time, while improving performance.

With an 8 Gbps FICON environment, you must address the following requirements:

- ▶ Storage growth: Number of ports and bandwidth
- ▶ Fabric growth: Number of Inter-Switch Links (ISLs) and scalability
- ▶ Performance: Bandwidth and cabling infrastructure
- ▶ Server virtualization: I/O throughput and scalability
- ▶ Services: End-to-end management
- ▶ Backup: Bandwidth
- ▶ Operational flexibility: Scalability and efficiency
- ▶ Compatibility: End-to-end management and lifetime

There are two System z-qualified FICON Director platforms that support 8 Gbps: the IBM System Storage b-type family and the Cisco MDS 9500 Series. Chapter 3, “FICON Director technical description” on page 45, explains the differences in proprietary implementations for each product line.

Some benefits of using 8 Gbps FICON are:

- ▶ The ISL oversubscription ratio is halved by upgrading from 4 Gbps to 16 Gbps ISLs.
- ▶ At long distances, 8 Gbps compared to lower speeds optimizes the number of dark fiber or WDM links.

- ▶ Building a high-performance infrastructure that provides the flexibility to deploy 8 Gbps as needed simplifies the fabric management and scales with the growth of applications and data.

We recommend that you do not mix low speed and high speed devices on the same 8 Gbps channel. The physical layer up to 4 Gbps link data rate had different limitations compared to 8 Gbps and 10 Gbps. Also, the 8 Gbps features only auto-negotiate to speeds of either 8 Gbps, 4 Gbps, or 2 Gbps.

We also recommend that you review the capabilities of the existing cabling infrastructure.

For more detailed information, refer to:

- ▶ Chapter 2, “System z FICON technical description” on page 21
- ▶ Chapter 3, “FICON Director technical description” on page 45
- ▶ *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *IBM System Storage SAN768B-2*, TSD03037USEN
- ▶ *Cisco MDS 9506 for IBM System Storage*, TSD00069USEN
- ▶ *Cisco MDS 9509 for IBM System Storage*, TSD00070USEN
- ▶ *Cisco MDS 9513 for IBM System Storage*, TSD01754USEN

### 4.4.3 Migrating from a single site to a multi-site environment

Migrating from a single site to a multi-site configuration usually means going from a point-to-point or switched topology to a cascaded topology.

Cascaded support is important for disaster recovery and business continuity solutions. It can provide high-availability connectivity and the potential for fiber infrastructure cost savings for extended storage networks. Solutions, such as IBM Geographically Dispersed Parallel Sysplex™ (GDPS), can benefit from the reduced intersite configuration complexity that FICON support of cascaded Directors provides.

Cascaded FICON Directors allow for FICON channels to connect a System z server to another, similar server or to peripheral devices (such as disk, tape library, or printer) using two FICON Directors.

In cascaded environments, the cabling for both intersite and intrasite can be simplified, which decreases the management complexity. It also reduces the number of FICON channels, Director ports, and storage devices. Sharing links between the two sites reduces the number of physical channels between sites and the number of Director ports. The faster the channel speeds between sites, the better the intersite throughput. With higher speed this factor can even multiply.

FICON Directors typically have more buffer credits per port than the server and the disk or the tape subsystems. The more buffer credits you have, the greater distances you can go without impacting the response times.

The following factors affect the performance of a cascaded FICON Director configuration:

- ▶ Number of ISLs and the routing policy
- ▶ Number of FICON channels routed across the ISLs
- ▶ ISL bandwidth
- ▶ ISL traffic management
- ▶ I/O workload characteristics
- ▶ Distances between the components
- ▶ Buffer-to-buffer credits

FICON support of cascaded Directors (sometimes referred to as cascaded switching or two-switch cascaded fabric) is for single-vendor fabrics only.

#### 4.4.4 Implementing a new FICON environment

When you are defining a new FICON environment, we recommend that you take advantage of the latest FICON functions and features that support high performance and high availability. For planning this scenario you must also consider criteria, such as:

- ▶ Facility requirements: Floor space, power, and cooling
- ▶ Scalability and flexibility: Ability to add ports and bandwidth dynamically
- ▶ Performance: I/O throughput and traffic patterns
- ▶ Availability: Redundant ports, bandwidth, and logical paths
- ▶ End-to-end management
- ▶ Cabling infrastructure
- ▶ Distances between the components
- ▶ Control unit and device capabilities
- ▶ Interoperability of all components

The remaining sections in this chapter help you to define a baseline for your FICON infrastructure, using best practices and industry standards that also align with vendor recommendations.

### 4.5 Topologies and supported distances

The unrepeated distance supported by System z FICON features is dependent upon:

- ▶ Transceiver type
- ▶ Fiber optic cabling type
- ▶ Port speed
- ▶ Mode conditioning patch (MCP) cables

Cabling specifications are defined by the Fibre Channel - Physical Interface - 4 (FC-PI-4) standard and used by System z FICON features. Table 4-1 identifies cabling types and link data rates that are supported in the FICON environment, including their allowable maximum distances and link loss budget.

The link loss budget is derived from the channel insertion loss budget defined by the FC-PI-4 standard (Revision 8.00).

Table 4-1 Fiber optic cabling for FICON - maximum distances and link loss budget

| FC-PI-4                                | 1 Gbps                       |                        | 2 Gbps                       |                        | 4 Gbps                       |                        | 8 Gbps                       |                        | 10 Gbps ISL <sup>a</sup>     |                        |
|--|------------------------------|------------------------|------------------------------|------------------------|------------------------------|------------------------|------------------------------|------------------------|------------------------------|------------------------|
|  | Distance in meters (in feet) | Link loss budget in dB | Distance in meters (in feet) | Link loss budget in dB | Distance in meters (in feet) | Link loss budget in dB | Distance in meters (in feet) | Link loss budget in dB | Distance in meters (in feet) | Link loss budget in dB |
| Fiber core (light source) <sup>b</sup> |                              |                        |                              |                        |                              |                        |                              |                        |                              |                        |
| 9 μm SM (10 km LX laser)               | 10000 (32736)                | 7.8                    | 10000 (32736)                | 7.8                    | 10000 (32736)                | 7.8                    | 10000 (32736)                | 6.4                    | 10000 (32736)                | 6.0                    |
| 9 μm SM (4 km LX laser)                | 4000 (13200)                 | 4.8                    | 4000 (13200)                 | 4.8                    | 4000 (13200)                 | 4.8                    | NA                           | NA                     | NA                           | NA                     |
| 50 μm MM <sup>c</sup> (SX laser)       | 860 (2822)                   | 4.62                   | 500 (1640)                   | 3.31                   | 380 (1247)                   | 2.88                   | 150 (492)                    | 2.04                   | 300 (984)                    | 2.6                    |

| FC-PI-4                               | 1 Gbps        |      | 2 Gbps       |      | 4 Gbps       |      | 8 Gbps      |      | 10 Gbps ISL <sup>a</sup> |     |
|---------------------------------------|---------------|------|--------------|------|--------------|------|-------------|------|--------------------------|-----|
| 50 μm MM <sup>d</sup><br>(SX laser)   | 500<br>(1640) | 3.85 | 300<br>(984) | 2.62 | 150<br>(492) | 2.06 | 50<br>(164) | 1.68 | 82<br>(269)              | 2.3 |
| 62.5 μm MM <sup>e</sup><br>(SX laser) | 300<br>(984)  | 3.0  | 150<br>(492) | 2.1  | 70<br>(230)  | 1.78 | 21<br>(69)  | 1.58 | 33<br>(108)              | 2.4 |

- Inter-Switch Link (ISL) between two FICON Directors
- In a single mode jumper cable, the minimum distance between connectors or splices is 4 meters (13.1 ft)
- OM3: 50/125 μm laser optimized multimode fiber with a minimum overfilled launch bandwidth of 1500 MHz-km at 850nm and an effective laser launch bandwidth of 2000 MHz-km at 850 nm in accordance with IEC 60793-2-10 Type A1a.2 fiber
- OM2: 50/125 μm multimode fiber with a bandwidth of 500 MHzkm at 850 nm and 500 MHz-km at 1300 nm in accordance with IEC 60793-2-10 Type A1a.1 fiber
- OM1: 62.5/125 μm multimode fiber with a minimum overfilled launch bandwidth of 200 MHzkm at 850 nm and 500 MHz-km at 1300 nm in accordance with IEC 60793-2-10 Type A1b fiber

**Note:** IBM does not support a mix of 50 μm and 62.5 μm fiber optic cabling in the same physical link.

Refer to 2.2, “System z FICON feature support” on page 37 for details about the FICON features available on the System z196, z114 and z10 servers.

## 4.5.1 Point-to-point

As illustrated in Figure 4-1, the maximum unrepeated distance of a FICON LX link using a single mode fiber optic cable is:

- ▶ 10 km (6.2 miles) for 1, 2, 4, 8 Gbps LX links
- ▶ 20 km (12.43 miles) for 1 Gbps LX links with RPQ 8P2263
- ▶ 12 km (7.46 miles) for 2 Gbps LX links with RPQ 8P2263

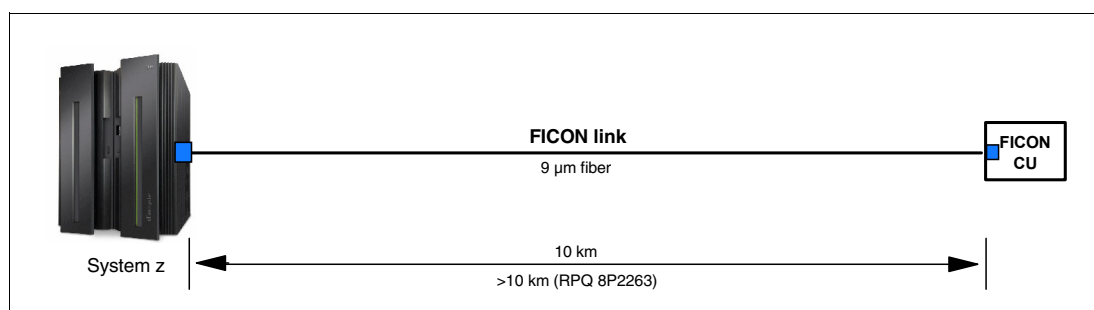


Figure 4-1 FICON LX point-to-point maximum unrepeated distance

## 4.5.2 Switched

As illustrated in Figure 4-2 on page 87, one FICON Director increases the maximum supported distance of a FICON LX channel path using single mode fiber optic cables to the following distances:

- ▶ 20 km (12.43 miles) for 1, 2, 4, 8 Gbps LX links
- ▶ 30 km (18.64 miles) for 1 Gbps LX links with RPQ 8P2263
- ▶ 22 km (13.67 miles) for 2 Gbps LX links with RPQ 8P2263




Figure 4-2 FICON LX switched maximum unrepeated distance

### 4.5.3 Cascaded

A FICON channel path can include a maximum of two cascaded FICON Directors. Using extended distance longwave transceivers on the Inter-switch links (ISL) between the FICON Directors might be required. Each ISL requires one fiber trunk (two fibers) between the FICON Directors.

As illustrated in Figure 4-3, and using the example of 10 km ISL links between FICON Directors, the maximum supported distance of a FICON LX channel path using single mode fiber optic cables is:

- ▶ 45 km (18.64 miles) for 1, 2, 4, 8 Gbps LX links
- ▶ 40 km (24.86 miles) for 1 Gbps LX links with RPQ 8P2263
- ▶ 32 km (19.88 miles) for 2 Gbps LX links with RPQ 8P2263
- ▶ 25KM (15.53 miles) for 8 Gbps LX ISL links




Figure 4-3 FICON LX cascaded FICON Directors maximum unrepeated distance

A FICON channel path through one or two FICON Directors consists of multiple optical fiber links. Each link in the channel path can be either long wavelength (LX) or short wavelength (SX), allowing the channel path to be made up of an intermix of link types as illustrated in Figure 4-4 on page 88.

This is possible because the FICON Director performs an optical-to-electrical-to-optical conversion (OEO) of the channel path as it passes through the Director. The transceiver types (LX or SX) at each end of a given link must match.




Figure 4-4 FICON Director media intermix

#### 4.5.4 Extended distance

The most important element for an extended distance topology is latency, which is usually caused by:

- ▶ The length of the fiber because of the time it takes for the light to travel over distance
- ▶ The opto-electro-opto converter (or the WDM equipment), if used

This must be compensated for by the buffer credit technology used inside the interface logic of the FICON Director or the channels at the System z server and the control unit.

The repeated distance for a FICON channel is System z-qualified to a maximum of 100 km (62 miles). For all FICON features using repeaters, the end-to-end distance between the FICON channel and the FICON Director port can be up to 100 km. The same end-to-end distance is also available between the FICON Director port and the control unit port. However, the overall end-to-end distance between the FICON channel and control unit is System z-qualified up to 100 km only.

**Note:** The maximum distance of the fiber optic link for a FICON Express8/8S LX feature to the first hop is limited to 10 km.

Using Dense Wavelength Division Multiplexing (DWDM) technology, the FICON SX and LX implementation supports a maximum distance of 100 km between two sites, as illustrated in Figure 4-5 on page 89.



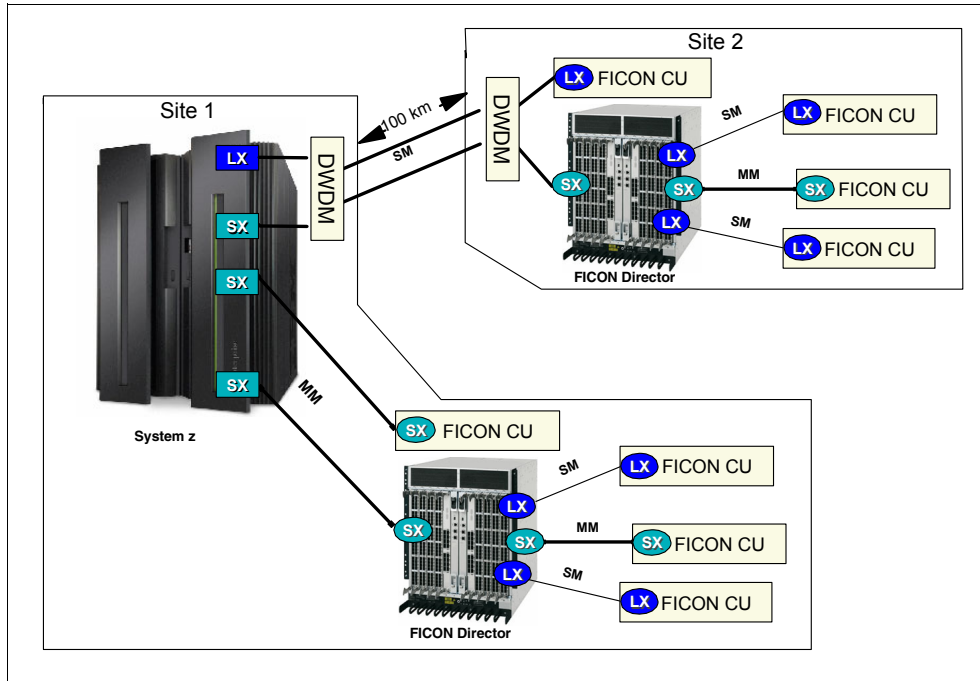


Figure 4-5 FICON point-to-point and switched point-to-point maximum extended distance

By combining FICON cascaded Director technology with DWDM technology, the FICON implementation supports a maximum distance of 100 km between two sites, as illustrated in Figure 4-6.

DWDM technology also provides increased flexibility because multiple links and protocol types can be transported over a single dark fiber trunk.

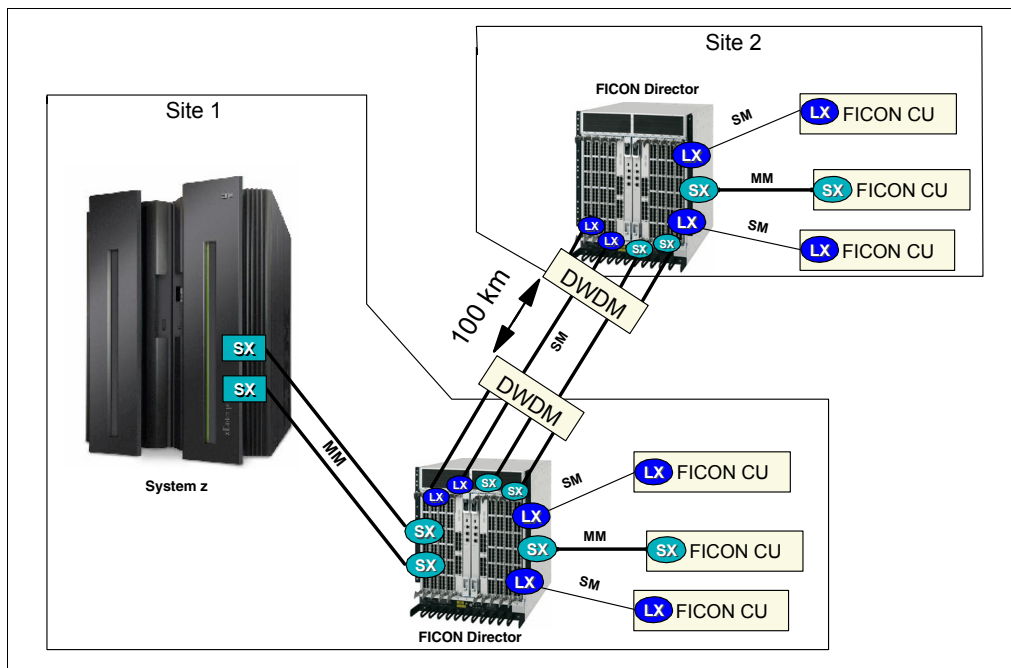


Figure 4-6 Cascaded FICON Directors maximum extended distance

For more detailed information, refer to:

*Planning for Fiber Optic Links (ESCON, FICON, Coupling Links, and Open Systems Adapters)*, GA23-0367

For the list of IBM qualified extended distance FICON Director solutions, see:

<http://www-03.ibm.com/systems/storage/san/index.html>

For the list of IBM qualified extended distance WDM solutions, see:

<http://www.ibm.com/servers/resourceLink>

## 4.6 Convergence

Some System z data center fabrics might require the use of a fully integrated multiprotocol infrastructure that supports mixed I/O and traffic types for simultaneous storage connectivity. A converged infrastructure enables you to consolidate multiple transport layers in a single physical interconnect. This consolidation provides the flexibility to build virtual server and storage environments supporting low-latency, high-bandwidth applications, and simplify hardware and cabling.

This approach also simplifies management and reduces power consumption, cooling, and space requirements inside a data center. Other relevant factors to consider are:

- ▶ High density port count (small form factor pluggable ports)
- ▶ Best energy effectiveness devices (per transferred byte)
- ▶ High availability (a redundant architecture of all components)
- ▶ High speed (switching capability - backplane throughput)

If you selected a point-to-point topology in the previous step, you can skip this section and go to 4.7, “Management” on page 94.

### 4.6.1 Intermix fabric

Here are reasons why you might move to an intermix FICON and FCP environment:

- ▶ Because of an open systems SAN that is delivering reliable, high-performance connectivity. These fabrics often have available ports on their SAN Directors.
- ▶ To consolidate switching and cabling infrastructure, which enables more effective maintenance, power consumption, cooling, and other operational tasks.

Possible scenarios include:

- ▶ FICON and open systems environments with a common storage network.
- ▶ System z servers accessing remote storage.
- ▶ Linux on System z to access local storage or remote storage.
- ▶ Hardware-based remote disk mirroring.
- ▶ Different System z server LPARs are accessing local or remote storage through FICON or using FCP.

Important technical considerations that you must take into account include fabric management techniques, such as zoning, partitioning, binding, and other design considerations. First you must consider the management techniques used in each protocol, and the management requirements of the target environment.

The management differences between the two protocols are not relevant unless you want to control the scope of the switching through zoning or connectivity control. For example, FCP devices use name server zoning as a way to provide fabric-wide connection control. FICON devices can use the Prohibit Dynamic Connectivity Mask (PDCM) to provide Director-wide connection control.

The PDCM is a vector-defined addressing system that establishes which addresses are allowed to communicate with each other. The FICON Director is required to support hardware enforcement of the connectivity control. In intermix environments, this is more restrictive than the zoning information used by open systems devices.

It is also important to understand the implications of FICON port addressing versus port numbering in FCP. FICON abstracts the concept of the port by creating an object known as the port address. An association is then made between the port address and the port number. This concept facilitates the ability to perform FICON port swaps for maintenance operations without the need to regenerate the host configuration. The port address abstraction is not in the Fibre Channel architecture and is foreign to FCP.

FCP communications are name-centric, discovery-oriented, and fabric-assigned, and they use the Fibre Channel Name Server to determine device communication. When an FCP device attaches to the fabric, it queries the Name Server for the zoning information. FICON devices do not query because the allowable port and device relationships have been previously defined.

FCP configurations support multiple fabrics and allow seven hops between source and target. FCP port-to-port connectivity has traditionally been enforced through zoning, although other techniques complementary to zoning (such as port, Director, and fabric binding) are used. Binding helps alleviate the security concerns that are experienced in intermix installations, because with FCP, any device in the fabric can access the ANSI standard FC management server by logging into the fabric.

If you are implementing intermix, you must block the transfer of any and all frames from a FICON Director port to all SAN connected ports and vice versa.

For more detailed information about these topics, refer to:

- ▶ *Implementing an IBM b-type SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *IBM/Cisco Multiprotocol Routing: An Introduction and Implementation*, SG24-7543
- ▶ *IBM System Storage b-type Multiprotocol Routing: An Introduction and Implementation*, SG24-7544
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545

## 4.6.2 Fabric security

Best security practices typically address three main objectives: availability and destruction, confidentiality and disclosure, and integrity and alteration of data.

However, it is important to remember that the majority of security threats are from insiders. Therefore, any SAN security strategy must also include:

- ▶ Restrict administrator privileges
- ▶ Isolate sensitive environments
- ▶ Audit SAN activities
- ▶ All-level documentation

Consider the following topics when planning a secure fabric.

- ▶ Securing access to the management tools, by using:
  - HTTPS
  - Secure Shell (SSH)
  - Secure Sockets Layer (SSL)
  - Secure Copy (SCP)
  - LDAP-MS Active Directory
  - Role-Based Access Control (RBAC)
  - Password policies
  - IP Filters (firewall)
  - Passive FTP
  - RADIUS server
- ▶ Securing the fabric, by using:
  - Zoning
  - Fabric Configuration Server (FCS)
  - Fabric Configuration Server (FCS) policy
  - Switch Connection Control (SCC) policy
  - Device Connection Control (DCC) policy
  - CHAP and DH-CHAP
  - IPsec
  - P-EAP and MS-CHAP
- ▶ Securing Long-Distance SAN, by using:
  - Physical separation
  - Electronic separation
  - Logical separation
- ▶ The use of Intrusion Detection and Incident Response:
  - Automatic event detection and management
  - Event forensics
  - Incident procedure to exercise due diligence
- ▶ The use of Fabric-Based Encryption

This technology enables you to encrypt specific or all disk data to prevent accidental data leaks when disk drives are replaced or disk arrays refreshed, and to protect enterprises from data theft. Although this is not supported with FICON, you can use the encryption capabilities of System z.

### 4.6.3 High integrity

High integrity ensures that any changes to the data streams are always detected, and that the data frames are delivered to the right endpoint.

A FICON channel requires the support of high-integrity architecture when two byte link addresses are configured on switched channels. During initialization, the FICON channel queries the Director to determine that it supports high-integrity architecture. If it does, the channel completes the initialization process, allowing the channel to operate with the fabric.

High-integrity fabric architecture support includes fabric binding and insistent Domain ID support.

A FICON Director provides the following fabric, Director, and port binding features:

- ▶ Fabric binding is a security method for restricting Directors within a multiple-Director fabric. This policy prevents unauthorized Directors from joining a fabric.
- ▶ Director binding is a security method for restricting devices that connect to a particular Director. A strict fabric-wide consistency is necessary for FICON Director binding.
- ▶ Port binding is a security method for restricting server or storage devices that connect to particular Director ports. This policy binds device ports to Director ports.

We recommend that you set the switch ID (using HCD/HCM) to the same value as the domain ID of the FICON Director, which is defined to the FICON Director at installation time. This simplifies the configuration and reduces confusion by having a common designation across all definitions. Domain ID and switch ID must be unique.

When defining the switch IDs in the IOCP definitions, ensure that you use values within the FICON Director's range.

If you use 2-byte addressing on the System z, strict fabric binding is mandatory (even with one FICON Director in the fabric).

## 4.6.4 Zoning

Zoning is a method used in the Director to restrict communication. There are two types of zoning: WWN zoning and port zoning.

Zoning is used, for example, to separate FICON devices from each other, to restrict confidential data to specific servers, and to control traffic paths.

The default zone automatically puts nodes attached to a fabric that are not already in a zone into a single zone and allow any-to-any connectivity. Conversely, default zone members are automatically removed whenever that member is added to an active zone.

The best practice recommendation is to disable the default zone and create port zones, even if the port zone includes all the ports in the fabric. By disabling the default zone and defining port zones, you ensure that the maximum number of zone members is not exceeded and does not inadvertently allow unintended connectivity.

WWN zoning permits connectivity between attached nodes based on WWN. That node can be moved anywhere in the fabric and it will remain in the same zone. WWN zoning is used in open systems environments and is *not* used for FICON. The IOCP effectively does what WWN zoning does in open systems environments. Adding WWN zoning on top of that adds an unnecessary layer of complexity because every time a new channel card or control unit interface is added to the fabric, the WWN zoning has to be modified.

Port zoning limits port connectivity based on ports. Use port zoning for FICON. Ports can easily be added to port zones, even if there is nothing attached to the port. Typically, all ports, with or without connections, in a FICON fabric are put into a single port zone to eliminate the need to make zoning changes when a new IODF is activated or channels and control units are added.

Only F-Ports (ports used for channel or control unit connections) need to be in a port zone. E-Ports (ports used for ISLs, connectivity between switches) do not need to be in a port zone. Although the recommended best practice is to not put E-Ports in the port zone, some

installations do so anyway so that ports can be re-allocated in the future without having to make zoning changes.

Define port zones for FICON and World Wide Name (WWN) zones for FCP. Port zones and WWN zones can be mixed in the same zone configuration but zones only have port members or WWN members.

Disk and tape mirroring protocols are FCP even when the front end is FICON. We recommend that you separate this traffic from the FICON traffic either with WWN zoning or virtual fabrics.

In more complex fabrics, more zones can provide additional security. In intermixed environments, only the FICON ports are put in the port zone.

FICON Prohibit Dynamic Connectivity Mask (PDCM) controls whether or not communication between a pair of ports in the Director is prohibited or allowed. If there are any differences in the restrictions set up and PDCM, the most restrictive rules are automatically applied.

Link incidents for FICON ports are reported only to registered FICON listener ports. The only exception to this is the loss of synchronization link incident.

To create a solid base for a SAN, use standard open systems best practices for creating WWN zones for the FCP traffic to ensure that the SAN is secure, stable, and easy to manage:

- ▶ Implement zoning, even if LUN masking is used.
- ▶ Persistently disable all unused ports to increase security.
- ▶ Use port WWN identification for all zoning configuration.
- ▶ Limit the aliases and names to allow maximum scaling.
- ▶ Use frame-based hardware enforcement.
- ▶ Use single initiator zoning with separate zones for tape and disk traffic if an HBA is used for both traffic types.
- ▶ Disable access between devices inside the default zone.

## 4.7 Management

We recommend that you manage the FICON fabric at various levels. Each level has a specific mission to fulfill. The most simple is offered by the command-line interface (CLI). Unfortunately, this does not provide enough control over the entire environment. To operate the FICON infrastructure you need to automate the whole process. All management levels create together robust systems management and interconnections with data center management.

### 4.7.1 Command-line interface

The command-line interface uses role-based access control (RBAC) to manage access for operations of the FICON Director. Changing the configuration of a FICON Director without a Telnet connection requires a serial connection. A serial connection might not be possible if the FICON Director is physically located in a restricted or remote area.

Because all FICON Directors ship with the same default IP address, you must change the IP address before connecting to the network. You also need a sufficient number of IP addresses for each FICON Director dependent on the platform. If all FICON Directors are connected to

the same management network, they must have the same subnet mask and have a common IP addressing range.

For recommended tasks executed from the CLI, refer to platform-specific documentation.

## 4.7.2 Element management

Web Tools enable an intuitive and easy-to-use interface that helps to monitor and manage FICON Directors. You can perform tasks by using a Java-capable web browser from your mobile computer, desktop PC, or workstation from any location with IP connectivity (Analogy to Cisco MDS Device Manager).

Web Tools provide an extensive set of features that enable you to quickly and easily perform key administrative tasks:

- ▶ Setting up IP addresses, Director names, and Simple Network Management Protocol
- ▶ Performing firmware upgrades and controlling boot parameters
- ▶ Maintaining RBAC, license keys, and RADIUS
- ▶ Enabling Ports on Demand
- ▶ Routing setup
- ▶ Configuring links and ISL trunks

## 4.7.3 Fabric management

One example of a client-server based application for managing the whole multi-fabric environment is the IBM Network Advisor.

IBM Network Advisor includes following:

- ▶ High scalability of managed Director ports
- ▶ Multiprotocol support
- ▶ Status and connectivity visualization
- ▶ End-to-end performance monitoring
- ▶ Fault management
- ▶ Troubleshooting and diagnostics
- ▶ Call home
- ▶ Offline support analysis
- ▶ Configuration management
- ▶ Security management
- ▶ FICON environment configuration and management with analysis and diagnostic capabilities
- ▶ Virtual fabric management

## 4.7.4 Storage management initiative specification

The storage management initiative specification (SMI-S) defines the interface that allows storage management systems to manage and monitor storage area network (SAN) resources. SMI-S is the only standard that addresses manageability from the perspective of a device in a storage array, all the way through the data path, to an application running on a host.

The SMI-S does not require any modification or upgrade to deployed fabrics when it is deployed.

The implementation is vendor-specific. It can be a daemon running directly on the Director platform or a separate host application. Each fabric device needs its own specific agent.

The advantage of using an end-to-end storage management platform is the possibility of management automation of the whole storage network including the FICON fabric.

For more details about IBM Storage Management Platform, refer to the following web site:

<http://www.storage.ibm.com>

## 4.7.5 System z management for FICON Directors

Combined, IBM Tivoli System Automation for z/OS and the control unit port (CUP) of the FICON Director can control and automate the entire I/O communication. IBM Resource Measurement Facility™ (RMF) and the CUP can provide activity statistics and resource utilization.

### Control unit port (optional feature)

The control unit port used by System z provides in-band management of the FICON Director. FICON management server mode must be enabled on the Director to enable CUP management features. This prevents local commands from interfering with host-based commands by starting serial access to Director parameters.

Use the CUP for the following reasons:

- ▶ It enables you to perform in-band management from the Hardware Management Console (HMC).
- ▶ It enables you to use the management tools used with ESCON, such as System Automation z/OS for I/O-Ops.
- ▶ It enables you to get Service Information Messages (SIMs) to the z/OS console for FICON device hardware failures.
- ▶ It enables you to use the Dynamic Channel Path Management function of Workload Manager and IRD (DCM) with FICON.
- ▶ It enables you to receive RMF Monitor I reports for reporting on FICON Director activity statistics in the SMF Record Type 74 Subtype 7. RMF Monitor I performs long-term data collection for resource utilization.

A FICON Director that supports CUP can be controlled by one or more host-based management tools. A mode register controls the behavior of the FICON Director according to CUP and other management interfaces. We make the following recommendations:

- ▶ Only one RMF should attempt to access the CUP at any given time. Too much activity to the CUP can cause missing interrupts.
- ▶ Best practice is to have two or more CHPIDs with access to the CUP for backup.

You also must configure the Missing Interrupt Handler process timeout, which determines when the channel will time out waiting for an expected response from the CUP. If processing times are increased, it is possible to inadvertently time out the channel because the CUP is busier than normal.

FICON Directors have an embedded port FE in the Control Processor for the CUP function. On a 256-port Director, this logical FE address overlaps the physical port FE address. When



using the CUP, on 256 port boxes, physical port addresses FE and FF cannot be used. In this case, use physical port addresses FE and FF for port swapping and for intermix ports. The following information applies to FICON Directors with at least 256 ports in use:

- ▶ The FICON Director is seen by the System z server as a 256-port Director.
- ▶ The Prohibit dynamic connectivity mask and port names are for ports 0 through 253. Ports 254 and 255 are reserved.

For information about the CUP license, behavior, and addressing rules, refer to FICON Director-specific documentation.

## **IBM Tivoli System Automation for z/OS**

With this management platform you can monitor, control, and automate a wide range of system elements spanning both the hardware and software resources of your infrastructure.

Tivoli System Automation for z/OS performs the following major operations:

- ▶ **System operations**  
It monitors and controls system operations, applications, and subsystems.
- ▶ **Processor operations**  
It monitors and controls processor hardware operations, and provides a connection from a focal point processor to a target processor.
- ▶ **I/O operations**  
It provides a single point-of-control for managing connectivity in your active I/O configurations. It takes an active role in detecting unusual I/O conditions and allows you to view and change paths between a processor and an input/output device, which can involve using dynamic switching, such as the FICON Director. I/O operations changes paths by allowing you to control channels, ports, Directors, control units, and input/output devices. You can do this through an operator console or API.

For more details, you can refer to:

*IBM Tivoli System Automation for z/OS Enterprise Automation, SG24-7308*

## **4.8 Virtualization and availability**

To achieve the most effective FICON infrastructure, consider virtualizing your FICON components. Virtualizing the FICON components enables a flexible and dynamic resource sharing solution. In addition, we always recommend that you have at least two physical connections to multiple FICON components to avoid single points of failures and ensure high availability.

In the following sections, we discuss the virtualization capabilities of the FICON components and what you need to take into account when considering high availability for your FICON infrastructure.

### **4.8.1 System z**

The use of FICON channels gives you the ability to define shared channels that can request logical paths. However, control units can only allocate a limited number of logical paths in relation to the number of logical paths that FICON channels can request. In configurations

where channels request more logical paths than a control unit can allocate, you must manage logical paths to help ensure that the I/O operations take place.

With proper planning, you can create I/O configuration definitions that allow control units in the configuration to allocate logical paths for every possible request made by channels in either of the following ways:

- ▶ Create a one-to-one correspondence between the logical path capacity of all control units in the physical configuration and the channels attempting to request them.
- ▶ Create I/O configurations that can exceed the logical path capacity of all or some of the control units in the physical configuration, but at the same time provide the capability to selectively establish logical connectivity between control units and channels as needed. This capability can be useful or even necessary in several configuration scenarios.

Several components of System z provide the capability for this virtual environment, as explained in the following sections.

### **Channel Subsystem**

Each server has its own Channel Subsystem (CSS). The CSS enables communication from server memory to peripherals through channel connections. The channels in the CSS permit transfer of data between main storage and I/O devices or other servers under the control of a channel program. The CSS allows channel I/O operations to continue independently of other operations within the server. This allows other functions to resume after an I/O operation has been initiated. The CSS also provides communication between logical partitions within a physical server using internal channels.

### **Multiple Channel Subsystem**

The design of System z servers offers a considerable increase in processing power, memory sizes, and I/O connectivity. In support of the larger I/O capability, the CSS has been scaled up correspondingly, and the Logical Channel Subsystem (LCSS) concept is designed to do that. Concepts were introduced to facilitate this architectural change and provide relief for the number of supported LPARs, channels, and devices available to the server.

### **Multiple Subchannel Sets**

Multiple Subchannel Sets (MSS) provides relief for I/O device configurations in large System z environments. It also increases Parallel Access Volume (PAV) connectivity. Subchannel numbers (including their implied path information to a device) are limited to four hexadecimal digits by hardware and software architectures. PAV has made this limitation of subchannels a challenge for larger installations. A single disk drive (with PAV) often consumes at least four subchannels. The solution gives you the ability to have sets of subchannels (addresses), with a current implementation of three sets on z196 and two sets on z114 and System z10.

For more information about this topic, refer to:

*Multiple Subchannel Sets: An Implementation View*, REDP-4387

### **Multiple Image Facility**

Multiple Image Facility (MIF) allows channels to be shared among multiple LPARs. Shared channels are configured to an LPAR, thus giving the LPAR a channel image of the shared channel that it can use. Each channel image allows an LPAR to independently access and control the shared channel as though it were a physical channel assigned to the LPAR.

By providing the logical equivalent of multiple physical channels dedicated to multiple LPARs, a shared channel can reduce hardware requirements without a corresponding reduction in I/O

connectivity. This reduction in hardware requirements can apply to physical channels, Director ports, and control unit ports, depending on the configuration.

MIF further improves control unit connection topologies for System z servers with multiple LPARs. MIF enables many LPARs to share a physical channel path, thereby reducing the number of channels and control unit interfaces required without a corresponding reduction in I/O connectivity.

Installation can take advantage of MIF performance enhancements offered by:

- ▶ Understanding and utilizing I/O-busy management enhancements
- ▶ Planning for concurrent data transfer
- ▶ Understanding examples of MIF consolidation
- ▶ Understanding and utilizing I/O-busy management enhancements

Before you can consolidate channels, you must be aware of the channel requirements of the particular control units you are configuring. The number of channels needed is independent of the number of LPARs on a system. The number of channels is based on the number of concurrent data transfers that the control unit is capable of. Although the recommended number of channels satisfies connectivity and performance requirements, additional channels can be added for availability.

Note that not all ESCON or FICON configurations benefit from the use of shared channels. There are some configurations where using an unshared channel is more appropriate, as explained here:

- ▶ When there are logical path limitations of the control unit  
Although many ESCON control units can communicate with multiple LPARs at a time using multiple logical paths, there are some ESCON-capable control units that can only communicate with one LPAR at a time.
- ▶ When the channel utilization of shared channels will be greater than unshared channels  
If you use shared channels to consolidate channel resources, you must consider the channel utilization of all the channels you consolidate. The channel utilization of a shared channel will roughly equal the sum of the channel utilizations of each unshared channel that it consolidates. If this total channel utilization can decrease performance, you must consider using unshared channels or a different configuration of shared and unshared channels to meet your connectivity needs.

**Note:** Current IBM tape and disk control units support FICON only

MIF allows you to use shared channels when defining shared devices. Using shared channels reduces the number of channels required, allows for increased channel utilization, and reduces the complexity of your IODF.

You cannot mix shared and unshared channel paths to the same control unit or device.

Channel spanning extends the MIF concept of sharing channels across logical partitions to sharing channels across logical partitions and channel subsystems. When defined that way, the channels can be transparently shared by any or all of the configured logical partitions, regardless of the Logical Channel Subsystem to which the logical partition is configured.

For more details see:

- ▶ *System z Enterprise Class System Overview*, SA22-1084
- ▶ *z/Architecture, Principles of Operation*, SA22-7832
- ▶ *System z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037

- ▶ *System z Processor Resource/Systems Manager Planning Guide*, SB10-7153
- ▶ *System z Support Element Operations Guide*, SC28-6096
- ▶ *IBM System z Connectivity Handbook*, SG24-5444

## 4.8.2 Control unit

Consider the following factors concerning the allocation of logical paths by control units:

- ▶ Control units allocate logical paths dynamically on a first come-first served basis. Control units do not manage the allocation of logical paths but instead allow channels to compete for logical paths until all of the control unit's logical paths are allocated.
- ▶ Control units vary in the number of logical paths they support.

The FICON infrastructure significantly affects the volume of logical path requests to a control unit as follows:

- ▶ Control units can attach to one or more ports on a Director or to additional ports on other Directors. Each Director port can dynamically connect to many other ports to which channels requesting logical paths are attached.
- ▶ For System z servers, each logical partition attaching to the same control unit competes for the control unit's logical paths.
- ▶ In a configuration where control units are shared by different servers, I/O configuration definitions for individual control units are not coordinated automatically among the IOCDs of the different servers. Each System z server competes for a control unit's logical paths.
- ▶ Shared channels require the establishment of a logical path for each channel image corresponding to an active LPAR sharing the channel. This can increase the number of logical paths that a single channel requests.

FICON support offers a high throughput performance for a single channel, which reduces the number of required channels. The CU channel adapters maintain compatibility with the existing infrastructure through auto-negotiation. A high performance FICON feature provides a significant reduction in channel utilization. This reduction can allow more I/O input on a single channel, and can also allow for a reduction in the number of FICON channels required to support a given workload. The number of offered CU channel adapters is further increasing the ability to share the storage.

On a DS8000 storage unit, each storage unit Fibre Channel adapter has four or eight ports. Each channel has a unique World Wide Port Name (WWPN). On FICON, the Fibre Channel CU channel adapter can operate with fabric or point-to-point topologies.

The CU FICON channels limit the number of devices per channel to 16384. To fully access 65280 devices on a storage unit, you need to connect a minimum of four FICON channels to the storage unit. You can access the devices through a FICON Director to a single storage unit FICON channel. With this method, you can expose 64 control-unit images (16384 devices) to each System z channel.

Parallel Access Volume (PAV) enables a single System z server to simultaneously process multiple I/O operations to the same logical volume, which can help to significantly reduce device queue delays (IOSQ time). This is achieved by defining multiple addresses per volume. With dynamic PAV, the assignment of addresses to volumes can be automatically managed to help the workload meet its performance objectives and reduce overall queuing. With PAV, reads are simultaneous. Writes to different domains (a set of tracks the disk controller is working on) are simultaneous as well. However, writes to the same domain are serialized. No double updates are possible to preserve integrity.

You can further enhance PAV by adding the IBM HyperPAV feature. IBM HyperPAV associates associate alias with base address on a per I/O basis. Therefore using fewer alias addresses. IBM HyperPAV is only supported on FICON channel paths.

For more information about these topics, refer to:

- ▶ *IBM System Storage DS8000 Architecture and Implementation*, SG24-8886
- ▶ *IBM System Storage DS8800 and DS8700 Introduction and Planning Guide*, GC27-2297

### 4.8.3 FICON Director

The Virtual Fabrics feature provides data, control, and management isolation for logical Directors and logical fabrics according the ANSI standard for Virtual Fabrics.

Virtualization allows you to:

- ▶ Isolate traffic types, applications, and management
- ▶ Pool resources
- ▶ Increase FICON consolidation and scalability

You must consider how many virtual fabrics for FICON with FICON Management Server you can create based on vendor limitations. All standard Director features are supported per virtual fabric. For a complete list of features refer to Chapter 3, “FICON Director technical description” on page 45 and to platform-specific documentation.

The Virtual Fabrics option can be added to the existing fabric capabilities, but it does *not* replace them. FICON cascading is only supported between switches from the same vendor.

The following topics influence the logical layout of virtual fabrics.

#### Port characteristics

The port addressing mode is platform-specific. For more details, refer to Chapter 3, “FICON Director technical description” on page 45 and to the platform-specific documentation.

For port addressing, use the worksheet provided in Appendix B, “Configuration worksheets” on page 405, to help document the physical layout of your FICON Director.

Ports automatically determine what they are connected to and at what speed they operate. Only use manual settings with equipment that cannot auto-negotiate or login properly. Typically, the only interfaces that require manual settings are older 1 Gbps interfaces and long distance ISLs.

Before installing the FICON Director, consider where you will connect your FICON channels, based on your availability requirements.

- ▶ Distribute the channels among different port cards.  
If two channels are defined to access the same CU, plug both fiber optic cables into different port cards in the Director.
- ▶ Distribute the CU ports among different port cards.  
If two paths are defined to attach the CU to the server through the Director, connect both fiber optic cables to ports on different cards.
- ▶ Distribute the ISLs across different port cards.  
If two or more ISLs are to be attached between the Directors, connect the fiber optic cables to ports on different cards. For frame-based trunking, the ISLs needs to be on the same ASIC.

Following these rules will ensure that there is always one path available between the server and the CU in case of a defective port card in the Director.

If multiple FICON Directors are planned based on the number of required connections and paths, consider spreading those paths across several Directors.

### **Inter-switch link**

The number of Inter-switch links must be carefully planned based on the existing and planned topology, workload characteristics, traffic patterns, and system performance. For availability, there must always be two ISLs between any two cascaded FICON Directors.

A general guideline is to plan for one ISL to four FICON channels that are routed over it. The ISL speed must reflect the speed of the channel.

We recommend that disk and tape must not share the same ISLs. Also FICON and FCP traffic must not share the same ISLs. Use TI zones to direct traffic to different ISLs. This is vendor-specific.

Plan on having enough ISL bandwidth so that the ISLs do not become congested when the maximum number of streaming tape devices are running because that can result in tape traffic termination.

### **Virtual ISL**

When two Directors are connected, an ISL connection for two logical switches or extended ISL (XISL) connection for two base switches is created, dependent on the configuration. When logical switches with the same Fabric ID are configured to use the XISL, the Director automatically creates a logical ISL (LISL) within the XISL. The LISL isolates traffic from multiple fabrics. XISL and LISL technology is not supported for FICON at this time.

For Cisco, the equivalent of virtual ISL is Trunking E-ports, which can transport virtual SAN Fabric information between Directors.

For information regarding the support of virtual ISLs in FICON environments, refer to platform-specific documentation.

## **4.9 Performance**

A single FICON channel can replace multiple ESCON channels. This means fewer channels, Director ports, and control unit ports, resulting in a simpler configuration to manage. Additionally, FICON supports the combined CTC and channel function.

ESCON-to-FICON channel aggregation can be anywhere between 8:1 and 2:1. An understanding of the factors that affect performance must be established before deciding on your target configuration. For example, the characteristics of your workload have a direct impact on the performance you realize in your FICON environment.

In addition, keep the utilization of your FICON channels at or below 50% utilization to maintain optimum response times.

The following items require analysis:

- ▶ I/O rates
- ▶ Block sizes
- ▶ Data chaining
- ▶ Read/write ratio

For more information about these topics, read the following performance papers:

- ▶ *Performance Considerations for a Cascaded FICON Director Environment Version 0.2x*, by Richard Basener and Catherine Cronin
- ▶ *IBM System z10 I/O and High Performance FICON for System z Channel Performance*, ZSW03059USEN

## I/O operations

Deploying FICON channels demands a further understanding of the I/O operations in your environment. The following terms describe the different phases with respect to the measurements available for determining the duration of an I/O operation.

- ▶ I/O supervisor queue time (IOSQ), measured by the operating system.

The I/O request might be queued in the operating system if the I/O device, represented by the UCB, is already being used by another I/O request from the same operating system image (UCB busy).

The I/O Supervisor (IOS) does not issue a start subchannel (SSCH) command to the Channel Subsystem until the current I/O operation to this device ends, thereby freeing the UCB for use by another I/O operation.
- ▶ Pending time (PEND), measured by the channel subsystem.

After IOS issues the start subchannel command, the channel subsystem might not be able to initiate the I/O operation if any path or device busy condition is encountered:

  - Channel busy
  - Director port busy
  - Control unit adapter busy
  - Device busy
- ▶ Connect time, measured by the channel subsystem.

This is the time that the channel is connected to the control unit, transferring data for the I/O operation.
- ▶ Disconnect time  

The channel is not being used for the I/O operation, because the control unit is disconnected from the channel, waiting for access to the data or to reconnect.

In the next sections, we describe factors that influence your FICON performance design.

### 4.9.1 Frame pacing

Frame pacing is an FC-4 application data exchange measurement and throttling mechanism. It uses buffer credits to provide a flow control mechanism for FICON to assure delivery of data across the FICON fabric.

After all buffer credits for a port are exhausted, a frame pacing delay occurs. These delays generally result in longer FICON connect time or longer PEND times that show up on the volumes attached to these links. *Frame latency* is the average amount of time it takes to deliver a frame from the source port to the destination port. Each element that handles the frame contributes to this latency measurement (channel, Director, storage port adapter, link distance, and so on). Only when using FICON Directors and when the CUP is enabled can RMF provide frame pacing delay information.

You must also keep in mind that tape workloads have larger payloads in a FICON frame, while disk workloads might have much smaller. The average payload size for disk is often about 800

to 1500 bytes. By using the FICON Director activity reports, you can gain an understanding of your average read and write frames sizes on a port basis.

The buffer credit represents the number of receive buffers supported by a port for receiving frames. The minimum value of buffer credits is one (1). This value is used as a controlling parameter in the flow of frames over the link to avoid possible overrun at the receiver.

The number of buffer credits is not associated with the performance until high data rates are attempted over long distances. If there are insufficient buffer credits, there might be a hard limit on the data rate that can be sustained.

The number of buffer credits available for each port on the FICON Director is implementation-dependent. The optimal amount of buffer credits is determined by the distance (frame delivery time), the processing time at the receiving port, link data rate, and the size of the frames being transmitted.

There are four implications to consider when planning buffer credit allocation:

- ▶ Ports do not negotiate buffer credits down to the lowest common value. A receiver simply advertises buffer credits to the linked transmitter.
- ▶ The exhaustion of buffer credits at any point between an initiator and a target limits overall performance.
- ▶ For write-intensive applications across an ISL (tape and disk replication), the buffer credit value advertised by the E\_Port on the target cascaded FICON Director is the major factor that limits performance.
- ▶ For read-intensive applications across an ISL (regular transactions), the buffer credit value advertised by the E\_Port on the local FICON Director is the major factor that limits performance.

Buffer credits are mainly a concern for extended distances; however, a poorly designed configuration can consume all available buffer credits and have an impact on performance. For example, assume you have a FICON 8 Gbps channel attached to two different control units running at lower link rates 4 Gbps and 2 Gbps. Depending on the traffic pattern, it can happen that the low speed device consumes all the available buffer credits of the 8 Gbps link, and you are not able to send more packets between the server and FICON Director. In this case, the average link utilization goes down. For this reason, we recommend that you upgrade all low speed channels where possible, or physically separate low and high speed environments.

For more information about this topic, refer to “Buffer credits” on page 51.

## 4.9.2 Extended distance FICON

The control units that support the Extended distance FICON feature can increase the IU pacing count, which eliminates handshakes between channel and control unit. This support allows the channel to remember the last pacing information and use this for subsequent operations to avoid performance degradation at the start of a new I/O operation. Improved IU pacing (255 IU instead of 16) improves the utilization of the link.

Extended distance FICON can eliminate the need for channel extenders in 2-site and 3-site z/OS Global Mirror configurations, because of the increased number of read commands simultaneously in flight. Extended distance FICON enables greater throughput over distance for IBM z/OS global mirror (XRC) using the same distance.



Note, however, that Extended distance FICON does *not* extend the achievable physical FICON distances or offer any performance enhancements in a non-z/OS global mirror environment.

For more information about this topic, refer to 2.1.11, “Extended distance FICON” on page 35.

### 4.9.3 Multiple allegiance

Multiple allegiance (MA) is the concurrent operations capability of the IBM disk storage systems that facilitates multiple accesses to or from the same volume with multiple channel path groups and system images. It allows multiple servers to establish concurrent implicit allegiances if there is no possibility that any of the channel programs can alter any data that another might read or write.

This can reduce the PEND time because device busy is not presented to the channel.

Resources that benefit from multiple allegiance are:

- ▶ Volumes that can have concurrent read operations or a high read-to-write ratio
- ▶ Data sets that have a high read-to-write ratio, multiple extents on one volume, or are concurrently shared by multiple users

### 4.9.4 Parallel Access Volume and HyperPAV

Parallel Access Volume (PAV) enables a single System z server to simultaneously process multiple I/O operations to the same logical volume, which can help to significantly reduce device queue delays (IOSQ time). This is achieved by defining multiple addresses per volume. With dynamic PAV, the assignment of addresses to volumes can be automatically managed to help the workload meet its performance objectives and reduce overall queuing. With PAV, reads are simultaneous. Writes to different domains (a set of tracks the disk controller is working on) are simultaneous too. However, writes to the same domain are serialized. No double updates are possible to preserve integrity.

HyperPAV is designed to:

- ▶ Provide a more efficient PAV function.
- ▶ For large volumes, to scale I/O rates without the need for additional PAV alias definitions.
- ▶ Exploit FICON architecture to reduce overhead, improve addressing efficiencies, and provide storage capacity and performance improvements:
  - Dynamic assignment of PAV aliases improves efficiency.
  - The number of PAV aliases needed might be reduced, taking fewer from the 64 K device limitation and leaving more storage for capacity use.
- ▶ Enable a more dynamic response to changing workloads.
- ▶ Simplified management of aliases.
- ▶ Allow migration to larger volume sizes.

HyperPAV provides a far more agile alias management algorithm than dynamic PAV because aliases are dynamically bound to a base for the duration of the I/O for the z/OS image that issued the I/O. When an I/O completes, the alias is returned to the pool for the logical control unit. It then becomes available to subsequent I/Os.

The number of aliases required can be approximated by the peak I/O rates multiplied by the average response time. For example, if the average response time is 4 ms and the peak I/O

rate is 2000 per second, then the average number of I/O operations executing at one time for that LCU during the peak is eight. Therefore, eight PAV-aliases must be all that is needed to handle the peak I/O rate for the LCU, along with all the other PAV-base addresses in the LCU.

Depending on the kind of workload, there is a huge reduction in PAV-alias UCBs with HyperPAV. The combination of HyperPAV and EAV allows you to significantly reduce the constraint on the 64 K device address limit and in turn increase the amount of addressable storage available on z/OS. In conjunction with Multiple Subchannel Sets (MSS) on System z, you have even more flexibility in device configuration.

For more information about this topic, refer to:

*Disk storage access with DB2 for z/OS*, REDP-4187

<http://www.redbooks.ibm.com/redpapers/pdfs/redp4187.pdf>

Use the PAV analysis tool to achieve a better understanding of your workload profile and achieve a more precise split:

[ftp://public.dhe.ibm.com/s390/zos/tools/pavanaanalysis/pav\\_analysis.doc](ftp://public.dhe.ibm.com/s390/zos/tools/pavanaanalysis/pav_analysis.doc)

## 4.9.5 Modified Indirect Data Address Word

Using Modified Indirect Data Address Words (MIDAWs) will not cause the bits to move any faster across the FICON link, but MIDAWs do reduce the number of frames and sequences flowing across the link, which makes the channel more efficient. Channel efficiency is essentially the relationship of throughput to channel utilization.

For a homogenous I/O stream of some uniform set of I/Os, the I/O rate, throughput, and channel utilization are all linearly related to each other. A channel becomes more efficient if the throughput at a particular level of utilization is increased.

The most significant performance benefit of MIDAWs is achieved with extended format data sets. The MIDAW facility is a conceptually simple modification to IDAWs. MIDAWs remove the 4 KB boundary restrictions of IDAWs. Using MIDAWs, the z/OS media manager can transfer a whole track using a single CCW. To span multiple tracks in one channel program, command chaining was used, and still is used. However, with the advent of MIDAWs, the media manager no longer uses data chaining.

If track-level CCWs were the foundation for improved FICON performance, MIDAWs represent the building that rests upon the foundation. MIDAWs and IDAWs have basically the same performance characteristics, and both perform better than data chaining for two reasons:

- ▶ One reason pertains to the number of CCWs used, which in turn affects the number of frames sent across the link. By reducing the number of frames, it takes less time for the FICON channel and the control unit host adapter port to process the channel program.
- ▶ The second reason pertains to FICON multiplexing and the fact that MIDAWs remove some of the data-chaining activity on the link. Multiplexing only affects performance when there are concurrent I/Os on the channel.

In addition to reducing channel utilization, indirect addressing also has the effect of reducing the utilization of the control unit FICON channel. This effect is hidden because the system does not have any way of measuring the utilization of the channel. The channel has to process every CCW, with or without data chaining, but indirect data addressing is transparent to the channel.

MIDAWs help reduce channel utilization and affect the relative performance of different block sizes. Physical disk performance is irrelevant if the data permanently resides in cache, or if the channels are faster than the speed at which data can be sequentially read from the disks.

However, block size can also affect FICON channel performance. In other words, in addition to eliminating the extended format performance penalty, using MIDAWs also eliminates the small block performance channel penalty, except to the extent that the record size affects the number of bytes that fit on a track. Small blocks are usually better for database transaction workloads, because the buffer hit ratio for random access tends to be higher with smaller blocks. MIDAWs affect the performance of update writes in the same manner as they affect reads.

For more information about this topic, refer to:

*How Does the MIDAW Facility Improve the Performance of FICON Channels Using DB2 and Other Workloads?*, REDP-4201

## 4.9.6 High Performance FICON

Substantial performance improvements can be realized in terms of throughput (MBps) and I/Os per second (number of starts per channel), depending on the application workload. The initial implementation of High Performance FICON (zHPF) by the IBM DS8000 is exclusively for I/Os that transfer less than a single track of data. Realistic production workloads with a mix of data transfer sizes can see between 30% and 70% of FICON I/Os utilizing zHPF. This can result in between a 10% to 30% saving in channel utilization.

The zHPF architecture has built-in reliability, availability, and serviceability enhancements. Significantly more information regarding errors is now provided by the control unit, thereby aiding fault isolation. There is also dedicated buffer for concurrent sense data now, thus making it unnecessary to separate sense commands. Missing interrupt problem determination is provided through an interrogate mechanism, thereby providing the operating system with the capability to query the state of operation at the control unit when detecting missing interrupts. Previously no information was available on missing interrupts and debugging involved attempting to correlate software, channel, and control unit log and trace data. Furthermore, the reason for aborts is provided to control units using an enhanced “purge path” mechanism.

There is no change required to the IODF when using the zHPF architecture. FICON channels are still defined the same way, as CHPID type FC. There is also no change required to the configuration or design of the fabric. I/O will flow over the fabric as either FICON or zHPF IUs. zHPF is transparent to FICON Directors.

The potential exists to achieve some channel aggregation with the adoption of zHPF, but remember to give careful consideration to meeting all your workload requirements. The benefits of zHPF particularly apply to OLTP I/O workloads such as DB2, VSAM, PDSE, and zFS which typically transfer small blocks of fixed (4 K) data.

zHPF-capable channels and devices support both FICON and zHPF traffic simultaneously. FICON continues to use CCWs and zHPF uses TCWs, so you have compatibility and coexistence.

For more information about this topic, refer to:

*High Performance FICON for System z: Technical Summary for Customer Planning*, ZSW03058USEN

## 4.9.7 Bandwidth management

This section discusses the features and function that can influence the fabric bandwidth.

### Local switching

There are two Director architectures:

- ▶ Multi-stage switching with interconnected ASICs  
The architecture approach uses ASICs to both process and switch data traffic. Processing can measure performance, enforcing zoning and other protocol-related tasks. Switching is simply directing traffic from one port to another. If more user-facing ports are required, two or more ASICs are internally interconnected.
- ▶ Single-stage switching with one or more crossbars  
ASICs process data on the port blades, but they do not switch the traffic. Switching is performed by one or more serial crossbars. This provides highly predictable performance across all ports and an easier design for higher port capacity. The bandwidth maximum for a Director chassis is the crossbar's bandwidth.

There are three connection allocation possibilities:

- ▶ You can gain a local switching benefit by distributing storage, server, and ISL ports evenly across Director port blades, to minimize the impact of losing a port blade to a hardware fault or other issue.
- ▶ You can perform a slight modification by distributing storage, server, and ISL ports evenly across port blade ASICs. By going across port blades and alternating between the high and low ASICs, local switching occurs much more frequently.
- ▶ You can group storage, server, and ISL ports that will need to switch with each other within Local Switching groups. Ensure that array devices are mapped to a local switching group that also has server ports that will require those devices.

Combinations of local and non-local traffic flows can occur simultaneously. The oversubscription rates can be optimized, depending on the design.

### Fabric Shortest Path First

Port speed-weighted Fabric Shortest Path First (FSPF) Fibre Channel routing protocol applies a weighting factor based on link speed to ensure that data traffic uses higher-speed paths between Directors. It gives a basic utilization of multiple ISLs between Directors with the same cost.

However, it does not assure that all links are equally utilized, which can cause a congestion on the link even others are not utilized at all. Because of this, we recommend that you use more effective ISL connectivity called frame-based trunking.

### Frame-based trunking

Using frame-based trunking, data flows are automatically distributed over multiple physical Inter-switch link connections. Frame-based trunking combines up to eight ISLs into a single logical trunk, optimizes link usage by evenly distributing traffic across all ISLs at the frame level, maintains in-order delivery to ensure data reliability, ensures reliability and availability even if a link in the trunk fails, and simplifies management by reducing the number of ISLs.

Frame-based trunking is the most effective way to utilize ISLs for FICON traffic while maintaining in-order-delivery (IOD) of frames inside the fabric.

## Exchange Based Routing

With Exchange Based Routing (EBR), also known as Dynamic Path Selection (DPS), the routing path is based on the Source ID (SID), Destination ID (DID), and Fibre Channel originator exchange ID (OXID). Every exchange can take a different path through the fabric. EBR requires the use of the Dynamic Load Sharing (DLS) feature. When there are multiple paths to a destination, the input traffic will be distributed across the different paths according to the bandwidth available on each of the paths. Every time there is a change in the network, the input is redistributed across the available paths. EBR is currently not supported for FICON, but can be utilized in a different logical switch on the same chassis when the virtual fabric feature is enabled.

## Port Based Routing

With Port Based Routing (PBR) the routing path is based only on the incoming port and destination port. Without the Lossless DLS feature enabled, after a route is established between a pair of ports that route remains until one of the port logs out. Routes are determined based on the Fabric Shortest Path First (FSPF) algorithm and available bandwidth at the time the route is created. In cascaded environment, this might result in unbalanced ISL utilization. When Lossless DLS is enabled, routes are re-calculated every time there is a fabric event resulting in better balancing of ISLs and internal paths for non-cascaded environments. The recommended best practice is to always enable Lossless DLS for FICON fabrics. The routing policies must be the same in all logical switches in the fabric but might be different on other logical switches on the same chassis.

## Lossless DLS

Lossless Dynamic Load Sharing (DLS) allows you to rebalance trunk port paths without causing input/output (I/O) failures in cases where the end devices require in-order-delivery (IOD) of frames.

## Port-channel

The Port-channel feature is an equivalent for link trunking on Cisco FICON Directors. The balancing is based either on a source-destination or source-destination-exchange ID policy. For FICON, only the source-destination policy is supported with guaranteed in-order-delivery of frames.

## Traffic isolation

The term “preferred path” for FICON traffic was replaced by the use of static routes and Traffic Isolation (TI) zones (not defined by the Fibre Channel standard).

Traffic Isolation allows data paths to be specified. It is used to:

- ▶ Separate disk and tape traffic
- ▶ Select traffic for diverse ISL routes
- ▶ Guarantee bandwidth for mission-critical data

Traffic isolation can operate with failover enabled or disabled.

Any traffic on ports that are not included in a TI zone follow the fabric shortest path first (FSPF) algorithm. TI zone can be used for failover traffic from other TI zones when the ISLs fail, and from ports that are not in a TI zone. It is recommended that you either put all ports in TI zones or do not use TI zoning at all.

Rules for traffic isolation zones (with IBM System Storage SAN b-type FICON Directors) are:

- ▶ Failover enabled is effectively a preferred path.
- ▶ Failover disabled creates an exclusive path.

- ▶ TI zones cannot be used to override the FSPF algorithm. If failover is enabled, ports not in a TI Zone will use the ISLs in the TI Zone when no other ISLs of equal cost exist. When failover is disabled, no path will be taken even if a higher cost path exists.
- ▶ Prior to FOS v6.4.0 a port only existed in one TI Zone.
- ▶ With FOS v6.4.0 and above a port can be in multiple TI Zones; however, there can only be one TI Zone to the same remote domain for a given port. In Figure 4-7, the TI Zone configuration is invalid because port 8 of Domain 3 is in two TI Zones, ETIZ 1 and ETIZ 2, but in Domain 1 port 1 is ETIZ 1 and port 4 is in ETIZ 2. To make this valid, port 1 or 4 of Domain 1 must move to a different domain.

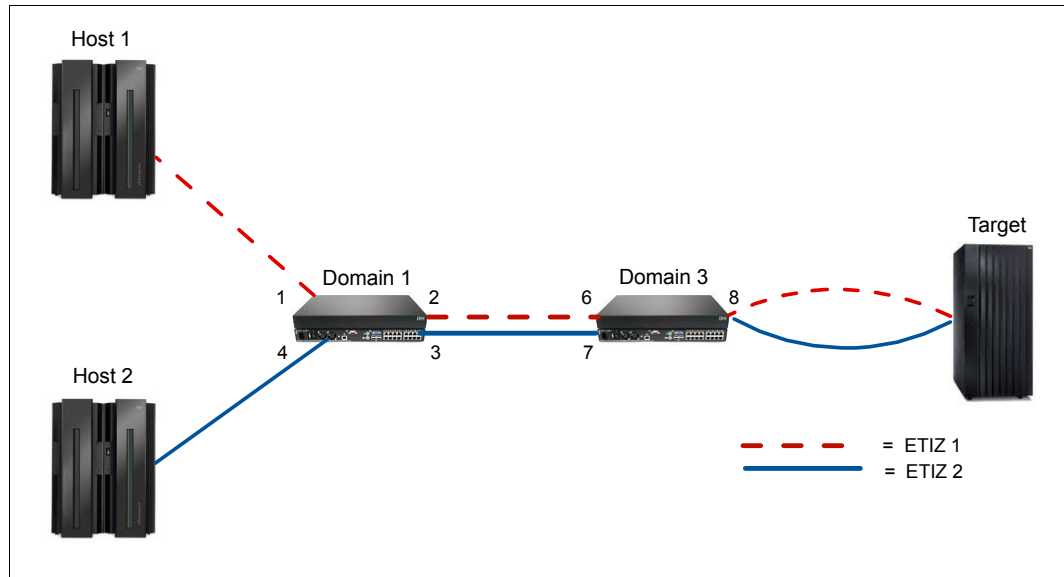


Figure 4-7 Invalid traffic isolation zones

## Port fencing

Bad optics can cause errors to occur at a high enough frequency that error processing and sending and processing RSCNs can cause fabric performance problems. Port fencing allows the user to limit the number of errors that a port can receive by forcing a port offline when certain error thresholds are met.

For FICON environments, port fencing is only set for CRC errors, Invalid Transmission Words, and Loss of Sync. The recommended thresholds for IBM System Storage SAN b-type FICON Directors are list in Table 4-2.

Table 4-2 Recommend Thresholds

| Parameter                 | Number of occurrences per minute |
|---------------------------|----------------------------------|
| C3 Discard Frames         | 2                                |
| Invalid Words             | 25                               |
| Invalid CRCs <sup>a</sup> | 3                                |
| Link Reset                | 2                                |
| Protocol Error            | 2                                |
| State Change <sup>b</sup> | 7                                |

- a. An ISL can travel through a path that is more subject to noise than internal data center connections so a slightly higher CRC threshold might make sense for E-Port connections in some configurations. In most cases the CRC is set to 3.
- b. The state change parameter is only available with FOS 6.4 and above.

These settings are high enough to ignore occasional errors and transient errors due to recabling, but low enough to stop problematic optics from causing fabric issues. By default, the alarms are set to fence the port, log an alert, send an email, and set an SNMP trap. In most FICON environments, only fencing the port and logging the alert are desired.

## 4.9.8 Traffic management

Virtual channel technology is the fundamental building block used to construct Adaptive Networking services. Virtual channels create multiple logical data paths across a single physical link or connection. They allocate their own resources (queues, flow control mechanisms such as buffer credits).

Adaptive Networking can assign a Quality of Service (QoS) priority to individual data flows in each virtual fabric. It anticipates congestion and dynamically makes adjustments in the fabric so that application traffic continues to flow.

Adaptive Networking introduces four new networking services: Quality of Service, traffic management, dynamic profiling, and resource recovery:

- ▶ QoS high, medium, and low priority are assigned to applications using zoning (Source ID/Destination ID pairs on WWN base).
- ▶ Automatic rate limiting and advanced queuing priority algorithms are implemented to remove congestion and dedicate bandwidth or assign data flow to specific path.
- ▶ Dynamic profiling provides knowledge about end-to-end resource use.
- ▶ Resource recovery prevents resource exhaustion and application disruption.

## 4.9.9 Evaluation tools

To maintain business-critical mainframe infrastructures you need the right tools, experience, and expertise. Without sufficient information, it is difficult to conduct capacity planning or predict the impact of FICON migration and other changes. For identifying performance bottlenecks, tracking the utilization levels of disk subsystems, channels, host adapters, and RAID ranks, you need additional performance data.

Using the suite of evaluation tools, you can provide the data and insight that are needed to effectively manage mainframe storage performance and optimize disk subsystem investments and fine-tune z/OS tape workloads.

The tools are listed and explained here:

- ▶ Processor evaluation: IBM zCP3000 study
- ▶ Channel evaluation: IBM FICON aggregation study

When planning to perform a capacity planning study for FICON channels, use the IBM System z Capacity Planning tool, zCP3000. zCP3000 is designed to support FICON channel aggregation (select channel candidates for aggregation onto FICON channels).

- ▶ Disk controller evaluation: IBM Disk Magic study

This is limited to the storage FICON channel port and port busy utilization numbers for workloads and configurations. You can model, if additional channels are needed.

- ▶ IBM zHPF evaluation study

This study is designed to quantify the amount of I/Os in an environment that appears to be zHPF-eligible. With this information, the impact/benefit of zHPF can be assessed for a specific situation or workload.

For System z, the z/OS Resource Measurement Facility (RMF) can be used to analyze your environment. This provides online, interactive performance monitoring or alternatively, long-term overview reporting with post-processor reports.

Some reports to assist in your analysis are:

- ▶ Channel path activity report

This report identifies each channel path by identifier and channel path type and reports on the channel utilization by server and individual LPAR.

- ▶ Device activity report

This report provides activity and performance information for selected devices.

- ▶ I/O queuing activity report

This report provides information about the I/O configuration and activity rate, queue lengths, and percentages when one or more I/O components were busy.

- ▶ FICON Director activity report

This report provides information about Director latency caused by port contention.

- ▶ Shared device activity report

This report gives you an overall performance picture of disk and tape devices that are shared between MVS systems in a sysplex.

For further information, refer to *z/OS RMF Report Analysis*, SC33-7991.

## 4.10 Prerequisites and interoperability

Device interoperability within a SAN is important because no single vendor can provide the entire infrastructure. In fact, most SAN environments have more than one server vendor and multiple storage vendors. SAN vendors provide comprehensive interoperability matrixes and have extensive experience designing, implementing, managing, and troubleshooting heterogeneous SANs.

Directors providing connectivity for storage and servers must also have common services and settings to operate with stability and full performance. When interconnecting Directors, the name server, zoning database, FSPF, RSCNs, and time-out value information must be shared.

Even if vendors follow Fibre Channel standards, coordinating fabric information between heterogeneous platforms is still an issue for fabric interoperability. Vendor implementation of the same standard can be at different stages and include varying enhancements, which can lead to fabric connectivity issues.

Native connectivity is required in the following situations:

- ▶ Migration between two platforms
- ▶ Port optimization within a fabric
- ▶ Temporary platform merge for infrastructure change
- ▶ Double vendor policies



Issues introduced by the native connectivity include:

- ▶ Management interface inconsistencies in fabric topology, node port states, zoning, and performance measurement
- ▶ Slow fabric recovery during adds, moves, and changes
- ▶ Failing fabric merges due to slight differences in standard implementations
- ▶ Name server synchronization issues
- ▶ Security features implementation
- ▶ Special vendor features to enhance fabric performance

E\_Port connectivity in these situations is used to form a single fabric. Using Fibre Channel Routing is an option to allow secure communication between hosts and storage in two or more separate, unmerged fabrics and to be managed independently. FCR is not supported by FICON. For this reason, any cascaded FICON Directors must be from the same vendor.

For System z details, refer to Chapter 2, “System z FICON technical description” on page 21.

For FICON Director details, refer to Chapter 3, “FICON Director technical description” on page 45.

- ▶ IBM System Storage web site:

<http://www-03.ibm.com/systems/storage/san/index.html>

- *IBM System Storage SAN768B Fiber Backbone Interoperability Matrix*
- *IBM System Storage SAN384B Fiber Backbone Interoperability Matrix*
- *Cisco MDS 9506, 9509, 9513 for IBM System Storage Directors Interoperability Matrix*

For more details regarding IBM System z qualification letters, see:

<http://www.ibm.com/servers/resourceLink>

## 4.11 Physical connectivity

The physical layer for FICON infrastructure is based on fiber optic cabling. When planning this infrastructure you need to consider the following items:

- ▶ Existing physical layer of the data center
- ▶ Existing FICON components
- ▶ Planned FICON components
- ▶ Non-FICON components that also require fiber cabling
- ▶ Capacity planning of the data center’s physical layer

The supported standard media for System z FICON features are single mode OS1 and OS2 and multimode OM1, OM2, and OM3 and OM4.

With the increasing data rate on the link you must utilize the latest fiber media types when designing the physical layer. New fiber media types provide more flexibility and better performance. A unified media platform is also important inside the data center for simplified operations and troubleshooting.

The highest flexibility and investment protection is with the single mode fiber. We recommend using the OS2 fiber type for System z FICON LX features because of the lower attenuation compared to the OS1 fiber type.

The System z FICON SX features support multimode fiber. For the best performance and investment protection for multimode, we recommend using the OM3 or OM4 fiber type. There are higher bandwidth fiber cables on the market today, but their use can introduce a risk for interoperability and might not be supported by the equipment vendor.

**Note:** OM1 and OM2 fiber types are not recommended for laser-based systems.

When choosing the fiber media, we recommend using only qualified vendor fiber optic cabling components. Also, avoid mixing cabling components from different vendors even if they are all qualified.

Qualified cabling components (connectors and cables) follow a standard color code. They are delivered with a measurement protocol directly from the manufacturer, which guarantees a high quality cable. This is important because the higher the bit rate on the channel, the more stringent the requirements for the fiber and the connectors. That means low attenuation, high return loss, and high number of cycles for reuse.

## Data center cabling infrastructure

To allow for adding, moving, and changing devices inside a data center, multiple networks typically share a structured cabling infrastructure in a hierarchical star topology. Therefore, it is important to view the FICON interconnections as an integral part of the overall fiber cabling infrastructure. This means the passive cross-connect points and traces are shared. This type of approach also provides a highly scalable and flexible cabling infrastructure.

At the hub of the hierarchical star topology is the Main Distribution Area (MDA). In IBM terminology it is the Central Patching Location, and it contains a large quantity of patch panels, placed in cabinets. The MDA is the central distribution point in the data center for telecommunication networks, local area networks (LANs), and SANs. It is the central cross-connect of the data center for all fiber optic and copper cables.

The MDA uses backbone or horizontal cabling to extend connectivity to the Zone Distribution Area (ZDAs) in the server, storage, and Director environments of the data center. The use of high-density cabling between the ZDA and the MDA reduces congestion and improves manageability, scalability, modularity, and security in the data center infrastructure.

All passive cross-connects, local distribution points, and equipment outlets are based on high density patch panels with a unified connector type for better scalability. We recommend using device-specific components, for example MTP brackets designed for System z. This provides a better starting point for cable routing inside the device rack.

When planning the routing of cables inside the cabinets, remember to plan for a sufficient number of patch cords and length, according to the vertical and horizontal cable management inside the cabinets. Failure to do so can cause multiple outages during the operations of the devices, when moving or laying cables.

There are various types of high density connectors available that offer differing parameters such as robustness, ease of use, number of channels, and fiber link parameters. We recommend using a unified connector type for the entire infrastructure.

The standard-based high density connectors are LC, MU, and MTP. The IBM high-density connector type is referred to as the small form factor duplex fiber optic connector for the data center (SCDC). The SCDC was especially developed for the mainframe fiber optic cabling environment, starting with ESCON.

When designing a fiber infrastructure, also consider the polarity management. We recommend using the standard-based ANSI/TIA-568-B.1-7 Method C. This is the default for the IBM Fiber Transport System.

Each cabling infrastructure must be carefully documented. Consider the naming conventions for each environment specifically. When creating a naming convention for System z, keep in mind the format used in the FICON environment, for ease of use and understanding by the operators.

When planning the optical interfaces, pay extra attention to the power level alignment between the receiver and the transmitter on the channel between two components. Low power levels do not necessarily mean a bad cable, but rather the potential of port errors and resending of frames, which in turn impacts performance.

The FICON Express8/8S features utilize the existing single-mode and multimode cable plants. However, the 8 Gbps channel is more sensitive to the condition of the cable plant. The cable plant must satisfy the industry standard specification to minimize connector loss and reflections.

We strongly recommend you thoroughly analyze the fiber optic link to ensure that the cable plant specifications (total cable plant optical loss as well as connectors and splices return loss) are being met for that link length. The most common source of cable plant optical link problems is associated with the various connections in the optical link. Dust, dirt, oil, or defective connections can cause a problem with high speed channels such as 8 Gbps, although lower link data rates such as 2 Gbps, or 4 Gbps might not be affected.

If you are experiencing excessive bit errors on a link (regardless of the data link rate), we recommend that you first clean and reassemble the connections. Refer to *IBM Fiber Optic Cleaning Procedure, SY27-2604* for the procedure and materials required. The cleaning is best performed by skilled personnel. The cleaning procedure might need to be performed more than once to ensure that all dust, dirt, or oil is removed. For more detailed information about fiber optic link testing and measurement, refer to *System z Maintenance Information for Fiber Optic Links (ESCON, FICON, Coupling Links, and Open System Adapters), SY27-2597*.

For additional information, refer to the following URLs:

<http://www.tiaonline.org>

<http://www.cenelec.eu>

<http://www.ibm.com/services/siteandfacilities>

For an example of how you can use this decision-making process when planning a FICON environment, we provide a walkthrough of our FICON environment with a cascaded FICON Director topology in Appendix A, “Example: Planning workflow” on page 397.





## Part 3

# Configuring the FICON environment

This part provides you with a basic guide for configuring a FICON environment. We give step-by-step guidance with descriptions of the software components and definitions that are needed to make the FICON environment operational.





## Configuring a point-to-point topology

In this chapter, we describe all of the tasks that are required to configure and define a FICON environment for a point-to-point topology.

Establishing a point-to-point topology involves the following items:

- ▶ Description of our environment
- ▶ Tasks and checklist:
  - Tasks
  - Verification checklist
- ▶ Defining the new configuration:
  - Defining the channel, CU, and storage devices
  - Defining our configuration example using HCD
  - Defining our configuration example using HCM
  - How the definition looks as IOCP statements
  - Defining a FICON Channel-to-Channel (FCTC) connection
- ▶ Building and activating the IODF:
  - Building the production IODF
  - Steps to activate the production IODF
- ▶ Configuring the IBM Storage System DS8000:
  - Connecting the fiber optic cables
  - Configuring the channel paths and devices online
  - Verifying the installation

## 5.1 Establishing a point-to-point topology

This section describes the tasks that are required to configure a FICON point-to-point topology on a System zEnterprise 196 server, including configuring the attached DS8000 storage subsystem. Our example is based on an infrastructure that is already built and installed, including the physical installation of the following items:

- ▶ zEnterprise z196 server
- ▶ DS8000 storage subsystem
- ▶ Fiber cabling infrastructure

## 5.2 Description of our environment

Based on the considerations explained in Chapter 4, “Planning the FICON environment” on page 77, we are going to build a FICON point-to-point configuration that consists of an IBM System z196 server and an IBM DS8000 storage subsystem.

In our example, all the required infrastructure is already installed in our environment. The configuration is shown Figure 5-1.




Figure 5-1 FICON Point-to-point configuration

The z196 server (SCZP301) has one LPAR (A11) defined and activated. The system name running in this partition is SC30. The operating system running in this partition is z/OS V1R13. As a standard feature of the z196 servers, zHPF protocol is used for data transfer to the DS8000 storage subsystem.

The two FICON channels are defined as spanned across CSS0, CSS1, CSS2, and CSS3 (only CSS1 is shown in our example). The channels are FICON Express8S LX features (FC 0409), using PCHID 398 (CHPID 20) and PCHID 3E0 (CHPID 24).

An IBM System Storage DS8000 is connected to the z196 server using CHPID 20 and CHPID 24. The two host adapters installed in the DS8000 are longwave (LX) laser. Port



numbers 0203 and 0303 in the DS8000 are used to connect to the server. Two logical control units (9000 and 9100) are defined, which have devices 9000-90FF and 9100-91FF assigned. The DS8000 has the zHPF feature enabled.

A maximum unrepeated distance of 10km (6.2 miles) is supported by the longwave laser (LX) feature when using 9µm single mode (SM) fiber optic cable. The fiber optic cables have an LC Duplex connector at both ends to connect to the z196 FICON Express8S channels and to the DS8000 host adapters.

## 5.3 Tasks and checklist

This section provides an overview of required tasks and the checklist that we used to make sure all hardware and software prerequisites were met before starting the configuration tasks.

### 5.3.1 Tasks

Figure 5-2 shows the main steps required to define and activate a FICON point-to-point configuration.

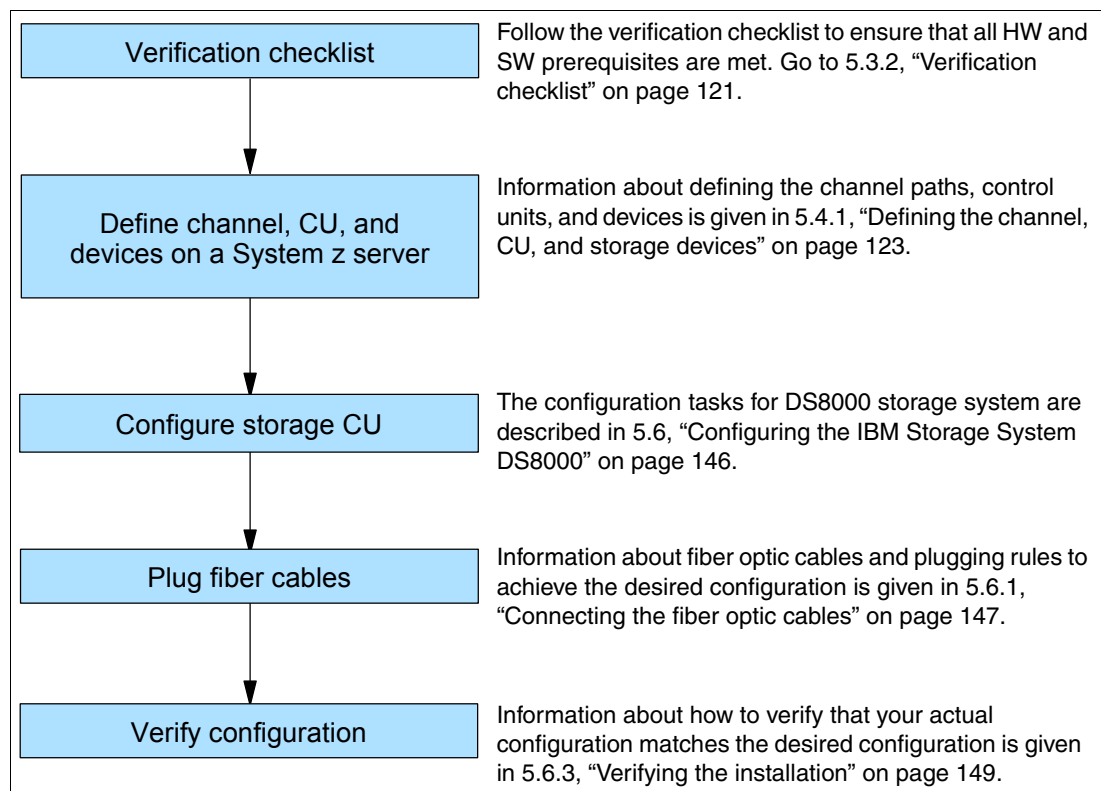


Figure 5-2 Steps for configuring and verifying a FICON point-to-point configuration

### 5.3.2 Verification checklist

Before configuring the point-to-point topology shown in Figure 5-1 on page 120, the following list was checked. All steps in the checklist must be finished and corrected (if required) to ensure a smooth and successful configuration of the topology. Both hardware and software requirements must be checked.

- Check that the appropriate FICON features are available on the System z server. For details about each feature, see “System z FICON feature support” on page 37.
  - FICON Express4
    - LX FC 3321, FC 3324
    - SX FC 3322
  - FICON Express8
    - LX FC 3325
    - SX FC 3326
  - FICON Express8S
    - LX FC 0409
    - SX FC 0410
- If using the FICON Express8S feature, check the System z operating system requirements.
  - Check the 2817DEVICE PSP bucket for the latest information about FICON Express8S support on the operating system.
- Check DS8000 storage hardware requirements to support FICON longwave (LX) or shortwave (SX) connectivity.
  - Check that FC 0709 and FC 7092 are installed to support zHPF.

**Note:** FC 0709 and FC 7092 must be ordered to obtain the license key to enable zHPF support in a DS8000 storage controller.

- Check that DS8000 firmware is at level 6.1.16.0 or higher to support the zHPF feature.
- Check that the zHPF feature is enabled in the DS8000.
- Check that the correct fiber optic cables are available to connect the z196 server to the DS8000 storage controller.
  - A 9µm single mode (SM) fiber optic cable is required to support longwave laser (LX) for a maximum distance of 10km (6.2 miles).
  - A 50µm or 62.5µm multi-mode (MM) fiber optic cable is required to support shortwave laser (SX). See “System z FICON feature support” on page 37 for the maximum supported distance depending on cable type and speed.
  - An LC duplex connector is required at both ends of the fiber optic cables to connect to the System z server FICON adapter and the DS8000 host adapter.

**Note:** All fiber optic cables used in a link must be of the same type. For example, they must be either all single mode or all multi-mode fiber optic cables.

## 5.4 Defining the new configuration

This section describes all the tasks that are required to reach the designated configuration shown in Figure 5-1 on page 120. It explains what must be done to get the FICON channels, storage control units, and the storage devices online and operating.

For our scenario, we had an active partition in a zEnterprise 196 server running z/OS V1.13 with HCD. We used HCD to create, save, and activate our I/O definitions.

We also show the definition process using HCM.

### 5.4.1 Defining the channel, CU, and storage devices

There are several tools available that you can use to define and configure a FICON environment on a System z server. The definition tools and an explanation of the statements and keywords used in the IOCP are described in Appendix C, “Configuration and definition tools” on page 409. Only statements and keywords relevant to our desired configuration are explained. Refer to *IOCP User’s Guide*, SB10-7037, for details about all available IOCP statements and keywords.

Steps using both HCD and HCM illustrate the definition process for the CHPIDs, CUs, and devices for the test configuration we describe in this section. Refer to *HCD User’s Guide*, SC33-7988 or *HCM User’s Guide*, SC33-7989, for more information and procedural examples.

After all definition tasks are completed, the definition data is stored in the I/O definition file (IODF) data set. Any time you make changes to the I/O configuration in the “Work” IODF, a new “Production” IODF must be built and saved to the IODF volume (software) along with updating and writing to a spare IOCDS (hardware). This is required to maintain synchronization of the IODF between the software and hardware definitions.

The newly created “Production” IODF is used to write IOCDS information to the Support Element of the System z server. The Support Element in the System z server stores the IOCDS on its built-in HDD. The IOCDS stored in the Support Element is loaded into the server’s storage (HSA) during a Power-on Reset (POR) of the System z server.

Changes or additions of channel paths, CUs, or devices can be performed dynamically in most cases. Because no system outage is required for dynamic changes, this is the preferred method for changing I/O definitions.

If changes to the I/O definitions are performed dynamically, those changes are directly updated into the Hardware System Area (HSA) of the server. The HSA is a separate storage area in the System z196 server that holds all I/O specifications and provides status information for all defined I/Os. The operating system uses the HSA to manage the I/O subsystem. For more details about System z196 channel subsystem characteristics, refer to 4.8.1, “System z” on page 97.

### 5.4.2 Defining our configuration example using HCD

To define our configuration example using HCD:

1. Select the current work IODF using HCD. We are using SYS6.IODF05.WORK for our example.
2. From the HCD main menu, select Option 1.3 Processor list.
3. Type *s* next to the processor ID to define/verify the CHPID definitions (for example SCZP301).
4. Type *s* next to the CSS ID, which contains the LPARs (for example A11 and A13).
5. Type *c* next to our CHPIDs to ‘Change Channel Path Definition’ (for example CHPID 20). Refer Figure 5-3 on page 124.

```

----- Change Channel Path Definition -----

Specify or revise the following values.

Processor ID . . . . . : SCZP301
Configuration mode . . : LPAR
Channel Subsystem ID : 1

Channel path ID . . . . . 20   +           PCHID . . . 398
Channel path type . . . FC     +
Operation mode . . . . . SPAN  +
Managed . . . . . No (Yes or No)   I/O Cluster _____ +
Description . . . . . Express8s LX

Specify the following values only if connected to a switch:

Dynamic entry switch ID ___ + (00 - FF)
Entry switch ID . . . . . ___ +
Entry port . . . . . ___ +

```

Figure 5-3 HCD Change Channel Path Definition

Here we can see that:

- Processor ID = SCZP301
- CSS = 1
- CHPID = 20
- PCHID = 398

6. Press Enter and observe the 'Define Access List' panel. Check that the required LPARs have access to CHPID 20. Refer Figure 5-4.

```

----- Define Access List -----
Row 19 of 43
Command ==> _____ Scroll ==> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 1
Channel path ID . . . : 20   Channel path type . . : FC
Operation mode . . . : SPAN   Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
/ 1   A11                1   OS   COMPLEX SC30
/ 1    A12                 2    OS    VMLINUX9
/ 1    A13                 3    OS    COMPLEX SC31
/ 1    A14                 4    OS
/ 1    A15                 5    OS
/ 1    A16                 6    OS    COMPLEX SC32
/ 1    A17                 7    OS    VMLINUX4
/ 1    A18                 8    OS    COMPLEX SC33
/ 1    A19                 9    OS    ITSOSI6
/ 2    A2A                 A    OS    ITSOSI3

```

Figure 5-4 HCD Define Access List

7. Perform the same steps for CHPID 24.
8. Return to the HCD main menu, and now select Option 1.4 Control units.
9. Scroll down the list until CU 9000 is located.
10. Type c next to CU 9000 to “Change Control Unit Definition”. See Figure 5-5.

```

----- Change Control Unit Definition -----

Specify or revise the following values.

Control unit number . . . . 9000 +
Control unit type . . . . . 2107          +
Serial number . . . . . 0L30002421 +
Description . . . . . YU CU 0

Connected to switches . . . 61 62 63 64  _  _  _  _  +
Ports . . . . . 1B 1B 10 10  _  _  _  _  +

Define more than eight ports . . 2  1. Yes
                                   2. No

```

Figure 5-5 HCD Change Control Unit Definition

In our example, this Control Unit already contains existing definitions. We are adding CHPIDs 20 and 24 to these definitions.

11. Press Enter to go to the Processor / CU panel.
12. Type 20 and 24 in the next available Channel Path ID. Link Address field. Refer Figure 5-6.

```

Select Processor / CU      Row 7 of 23 More:
Command ==> _____ Scroll ==> CSR

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

Control unit number . . : 9000      Control unit type . . . : 2107

-----Channel Path ID . Link Address + -----
/ Proc.CSSID 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
_ SCZP301.0 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____
_ SCZP301.1 40.1B  41.1B  42.10  43.10  20   24   _____ _____
_ SCZP301.2 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____
_ SCZP301.3 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____

```

Figure 5-6 HCD Adding CHPIDs 20 and 24 to SCZP301.1, CU 9000

13. Press Enter to observe the Modify Device Parameters panel. See Figure 5-7 on page 126.

```

Select Processor / CU
----- Modify Device Parameters -----
Row 1 of 3 More: >
Command ==> _____ Scroll ==> CSR

Specify or revise any changes to the device parameters in the list below.
To view attached control units, scroll to the right.

Processor ID . . . . : SCZP301
Channel Subsystem ID : 1

-----Device-----      ---UA----      Preferred Exposure
No., Range Type          SS+ Old New + Time-Out STADET CHPID + Device
9000,016 3390B          - 00 00 No Yes --
9000,128 3390A          1 80 80 No Yes --
9010,112 3390B          - 10 10 No Yes --
***** Bottom of data *****

```

Figure 5-7 HCD - Modify Device Parameters

This panel will display the Unit Addresses (UA) already defined for the device:

- Devices 9000-907F as 3390B (Subchannel Set or SS=0)
- Devices 9000-907F as 3390A (Subchannel Set or SS=1)

**Note:** With the addition of HyperPAV, we do not need to define 128 alias devices. Refer to 4.9.4, “Parallel Access Volume and HyperPAV” on page 105 to calculate your alias requirements.

14. Finally Return to the HCD main menu, and select Option 1.1 Operating system configurations.
15. Type u next to the Operating system for our example (for example TEST2094).
16. Scroll down the list until I/O Device 9000 is located.
17. Type c next to 9000, and press enter until the Define Device to Operating System Configuration panel is displayed.
18. Type s next to TEST2094 to start the Device Parameters / Features panel. Refer to Figure 5-8 on page 127.

```

----- Define Device to Operating System Configuration -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Select OSs to connect or disconnect devices, then press Enter.

Device number . : 9000          Number of devices : 16
Device type . . : 3390B

/ Config. ID  Type   SS Description          Defined
- ALLDEV     MVS   All devices           Yes
- LABSERV1   MVS   Lab Services          Yes
- LOGRMVS1   MVS   Sysplex systems      Yes
- MVSW1      MVS   Production systems   Yes
- OPENMVS1   MVS   OpenEdition MVS      Yes
s TEST2094   MVS   Sysplex systems      Yes

```

Figure 5-8 HCD - Defining a Device to an Operating System Configuration

Here we can enable/disable the device-specific features by over typing the yes/no values. Refer to Figure 5-9.

```

----- Define Device Parameters / Features -----
                                                    Row 1 of 6
Command ==> _____ Scroll ==> CSR

Specify or revise the values below.

Configuration ID . : TEST2094    Sysplex systems
Device number . . : 9000        Number of devices : 16
Device type . . . : 3390B

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 *****

```

Figure 5-9 HCD Define Device Parameters / Features

19. Finally, press Enter to go to the Assign/Unassign Device to Esoteric panel. Refer to Figure 5-10 on page 128.

```

----- Assign/Unassign Device to Esoteric -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Specify Yes to assign or No to unassign.  To view devices already
assigned to esoteric, select and press Enter.

Configuration ID : TEST2094           Sysplex systems
Device number   . : 9000             Number of devices  : 16
Device type    . . : 3390B          Generic . . . . . : 3390

+-oteric Assigned Starting Number Number of Devices
- 01.ATL3      No      _____ _____
- 01.CART      No      _____ _____
- 01.NOVIO     Yes    _____ _____
- 01.SYSDA     Yes    _____ _____
- 01.SYSSQ     No      _____ _____
- 01.TAPE      No      _____ _____
- 01.TAPE3592 No      _____ _____
- 01.VIO       Yes    _____ _____

```

Figure 5-10 HCD Assign/Unassign Device to Esoteric

This panel allows us to assign devices to particular esoterics, such as SYSDA and VIO.

20. Perform the same steps for the remainder of the devices for this control unit (3390B 9010-7F and 3390A 9000-7F), then return to the HCD main menu.

21. Steps 8 through 19 can now be performed for Control Unit 9100 and Devices 9100-917F.

### 5.4.3 Defining our configuration example using HCM

To define our configuration example using HCM:

1. Select the current work IODF using HCM. We are using SYS6.IODF05.WORK for our example. Refer to page Figure 5-11 on page 129.






Figure 5-11 HCM File Open IODF05.hcm

2. Select **Locate** → **Processor**.
3. Double-click the required processor in the list (for example SCZP301). See Figure 5-12.




Figure 5-12 HCM Processor display

Our configuration already contains existing definitions, so we will be adding to these.

4. Double-click the top pane of the processor display to edit the processor, or select **Edit** → **Processor**, and select the processor from the list. Refer to Figure 5-13 on page 130.

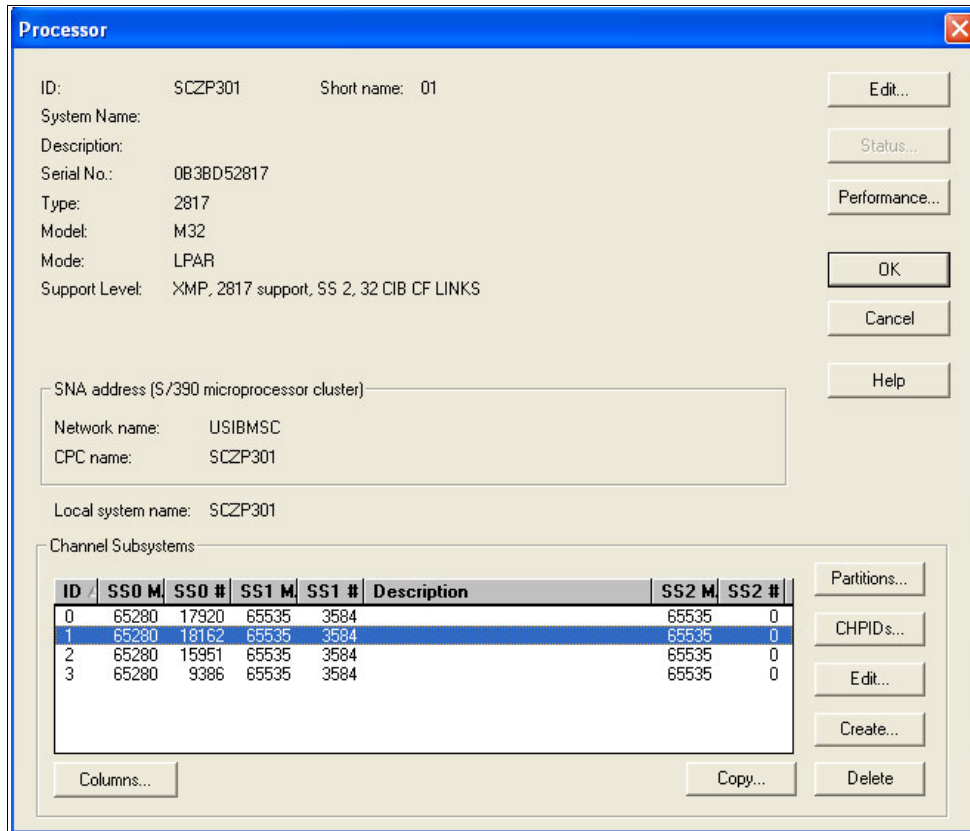


Figure 5-13 HCM Edit Processor

5. Highlight CSS ID 1, and click **CHPIDs**.
6. Scroll down the list until CHPIDs 20 and CHPIDs 24 are located. See Figure 5-14.

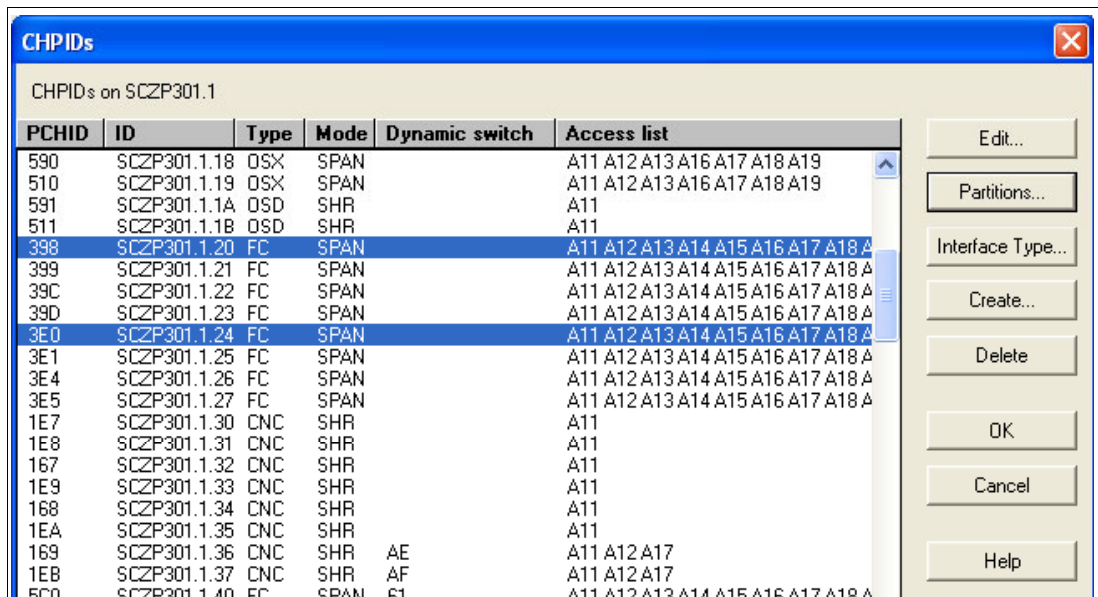


Figure 5-14 HCM Edit CHPIDs

7. Verify that they are defined correctly and have access to our test LPARs A11 and A13.

8. Select **Edit** → **Controller**, and select the DASD controller for our test. Refer to Figure 5-15.




Figure 5-15 HCM DASD Controller view

Currently our DASD subsystem has eight FICON connections spread over four FICON directors. For our example, we add two point-to-point (or direct connect) FICON connections to the DS8000.

9. Double-click the controller figure, and in the Interfaces box, highlight our two interfaces 0203 and 0303. Click **Connect**, as shown in Figure 5-16 on page 132.

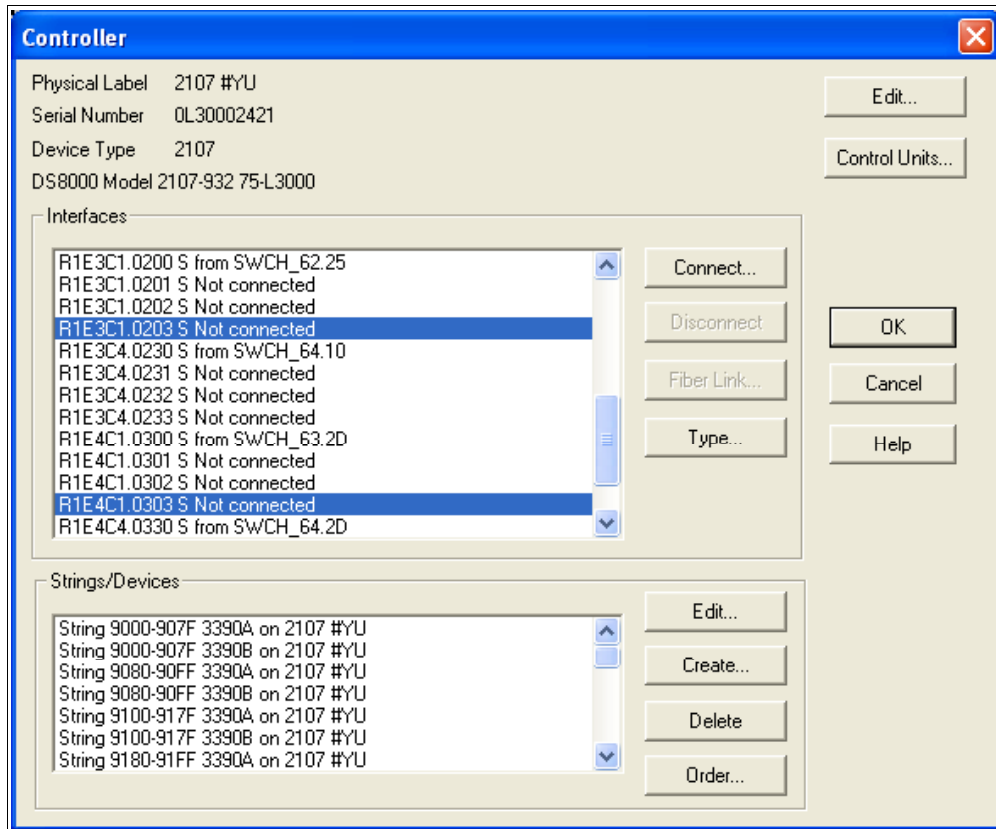


Figure 5-16 HCM Connecting to the DS8000 interfaces

10. Select CHPIDs 20 and 24 from processor SCZP301, and then click **OK**. See Figure 5-17.




Figure 5-17 HCM DASD point-to-point connections

Now we must ensure that the device definitions for the Operating System (software) are correct.

11. Double-click the controller box.

12. Click **Control Units**.

13. Highlight CU 9000. Click **I/O Devices**.
14. Highlight all of the 3390B devices (9000-907F). Click **OS Configs**.
15. Highlight TEST2094, and then click **Define**. Refer to Figure 5-18.

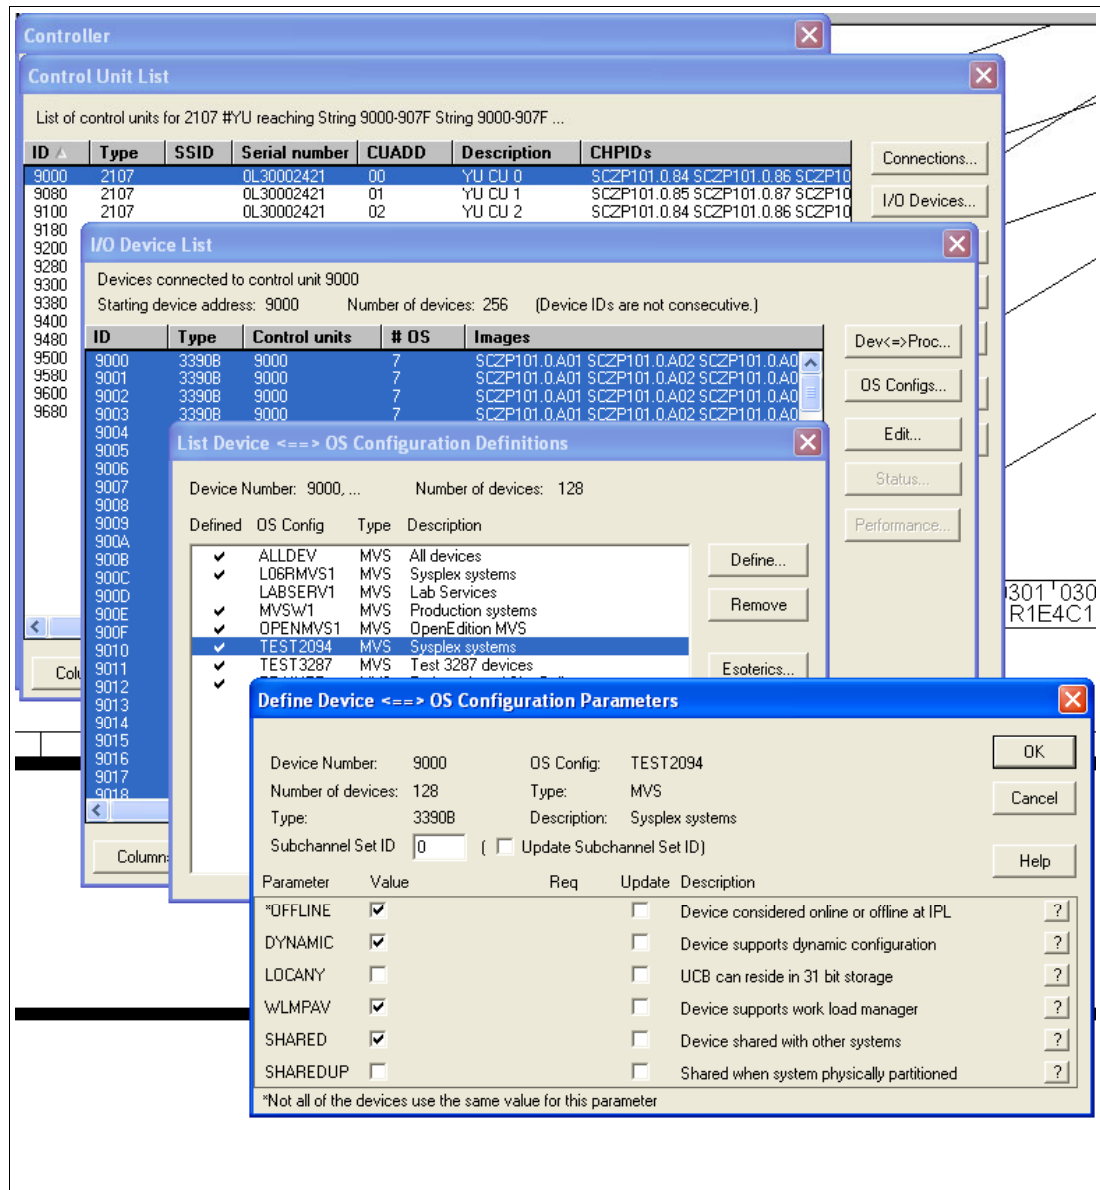


Figure 5-18 HCM OS Configuration Parameters

16. Verify that the Device Parameter values are correct.
17. Click **OK**, and then click **Esoterics**. Verify that the devices are defined to the correct Esoterics.
18. Perform the same operation for the 3390A devices for CU 9000 and for CU 9100 and Devices 9100-917F.

## 5.4.4 How the definition looks as IOCP statements

The following examples show the CHPID, CNTLUNIT, and IODEVICE statements when all of the definition tasks, based on the configuration shown in Example 5-1, are finished.

Example 5-1 shows the CHPID statement and its keywords.

*Example 5-1 CHPID definition for point-to-point configuration*

---

```

CHPID PATH=(CSS(0,1,2,3),20),SHARED, *
      PARTITION=((CSS(3),(A31,A34,A35),=)), *
      NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),=)),(CSS(1), *
      (A1B,A1E,A1F),=)),(CSS(2),(A2E,A2F),=)),PCHID=398, *
      TYPE=FC
CHPID PATH=(CSS(0,1,2,3),24),SHARED, *
      PARTITION=((CSS(3),(A31,A34,A35),=)), *
      NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),=)),(CSS(1), *
      (A1B,A1E,A1F),=)),(CSS(2),(A2E,A2F),=)),PCHID=3E0, *
      TYPE=FC

```

---

- ▶ The **PATH** keyword in the CHPID statement defines CHPID 20 and CHPID 24 in CSS(1) as SHARED.
- ▶ The **NOTPART** keyword specifies the LPARs which do not have access to CHPID 20 and CHPID 24.
- ▶ The **PCHID** keyword specifies the physical channel identification number associated with CHPID 20 and CHPID 24.

With the CHPIDs defined, we next show how the Control Units attached to the CHPIDs are defined. Example 5-2 displays the CNTLUNIT statement and keywords.

*Example 5-2 CNTLUNIT definition for point-to-point configuration*

---

```

CNTLUNIT CUNUMBR=9000, *
      PATH=((CSS(0),40,41,42,43),(CSS(1),40,41,42,43,20,24),(C *
      SS(2),40,41,42,43),(CSS(3),40,41,42,43)), *
      UNITADD=((00,256)), *
      LINK=((CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,**,**),(C *
      SS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)),CUADD=0, *
      UNIT=2107
CNTLUNIT CUNUMBR=9100, *
      PATH=((CSS(0),40,41,42,43),(CSS(1),40,41,42,43,20,24),(C *
      SS(2),40,41,42,43),(CSS(3),40,41,42,43)), *
      UNITADD=((00,256)), *
      LINK=((CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,**,**),(C *
      SS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)),CUADD=2, *
      UNIT=2107

```

---

- ▶ The **CUNUMBR** keyword shows definitions for 9000 and 9100.
- ▶ The **PATH** keyword shows CSS(1) as having definitions for CHPID 20 and CHPID 24.
- ▶ The **CUADD** keyword (representing the logical control unit or LCU) is set to 0 for CU 9000 and 2 for 9100. By pressing F20 in the panel showed in Figure 5-6 on page 125, you can see the **CUADD** value set.

After the CHPIDs and the CUs are defined, the next step is to define the devices owned by the CUs. Example 5-3 on page 135 displays the IODEVICE statement and keywords.

*Example 5-3 IODEVICE definition for point-to-point configuration*

---

```
IODEVICE ADDRESS=(9000,128),CUNUMBR=(9000),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9000,128),UNITADD=80,CUNUMBR=(9000),          *
      STADET=Y,SCHSET=1,UNIT=3390A
IODEVICE ADDRESS=(9100,128),CUNUMBR=(9100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9100,128),UNITADD=80,CUNUMBR=(9100),          *
      STADET=Y,SCHSET=1,UNIT=3390A
```

---

- ▶ The ADDRESS keyword shows device address 9000 for 128 (9000-907F) or 9100 for 128 (9100-917F).
- ▶ The CUNUMBR keyword associates the IODEVICE statement with the Control Unit number 9000 or 9100.
- ▶ The UNIT keyword shows the device type of 3390B (for the base devices) or 3390A (for the HyperPAV alias devices).
- ▶ The UNITADD keyword is used when the Unit Address (UA) does not start at 00.
- ▶ The SCHSET keyword shows that these devices are defined in Subchannel Set 1 (SS1).

**Note:** For CU 9000 and CU 9100, the device allocation is defined in the following way:

- ▶ 3390B devices 9000-907F, CUADD=0, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9000-907F, CUADD=0, UA = 80-FF, SCHSET = 1
- ▶ 3390B devices 9100-917F, CUADD=2, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9100-917F, CUADD=2, UA = 80-FF, SCHSET = 1

## 5.4.5 Defining a FICON Channel-to-Channel (FCTC) connection

In addition to storage control units and storage devices, we define some point-to-point FICON Channel-to-Channel (FCTC) CUs and devices to allow communication between LPARs on the same processor. For more information about considerations regarding FCTCs, refer to Appendix G, “Adding FICON CTC connections” on page 489.

Based on those considerations, we configure FCTCs in a point-to-point FICON environment. Figure 5-19 on page 136 illustrates the desired FCTC configuration. To simplify the illustration, only four CUs and two LPARs are shown in this sample configuration, although more than four CUs are normally defined.




Figure 5-19 FCTC configuration (point-to-point)

Both LPARs have access to CHPID 23 and CHPID 27 because the CHPIDs are defined as shared. There is a direct connection between CHPID 23 and CHPID 27 by a fiber optic cable.

Each LPAR has access to a logical control unit and its logical devices defined in the corresponding LPAR:

- ▶ LPAR A11 communicates using a FICON channel to CU 5134, which is defined to LPAR A13.
- ▶ LPAR A13 communicates through a FICON channel to CU 4114, which is defined to LPAR A11.

Example 5-4 shows the definition of FCTC control units and device definitions.

*Example 5-4 FCTC definitions point-to-point*

---

```

CNTLUNIT CUNUMBR=4114,PATH=((CSS(1),23)),UNITADD=((00,004)), *
      CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(4114,004),UNITADD=00,CUNUMBR=(4114), *
      STADET=Y,NOTPART=((CSS(1),A11,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=5134,PATH=((CSS(1),27)),UNITADD=((00,004)), *
      CUADD=13,UNIT=FCTC
IODEVICE ADDRESS=(5134,004),UNITADD=00,CUNUMBR=(5134), *
      STADET=Y,NOTPART=((CSS(1),A13,A1B,A1E,A1F)),UNIT=FCTC

```

---

- ▶ CUNUMBR keyword shows CU 4114 for LPAR A11 and CU 5134 for LPAR A13
- ▶ PATH keyword shows the CSS and CHPID that the CU is defined to
- ▶ CUADD keyword shows the logical control unit number (LCU) which must equal the LPAR ID
- ▶ NOTPART keyword shows the partitions that cannot access these devices. This is controlled using the Explicit Device Candidate List (EDCL) option in the Operating system configuration definition.



There are four devices defined for each LPAR (LPAR A11 and LPAR A13) to allow communications between both LPARs using their corresponding devices:

- ▶ LPAR SC30 (A11) communicates with LPAR SC31 (A13) using devices 5134-5137
- ▶ LPAR SC31 (A13) communicates with LPAR SC30 (A11) using devices 4114-4117
  - Devices 4114-4117 cannot be brought online to LPAR SC30 (A11)
  - Devices 5134-5137 cannot be brought online to LPAR SC31 (A13)

The rules we followed for device numbering are described in “FCTC device numbering scheme” on page 490.

Figure 5-20 illustrates the logical view of the FCTC configuration and the data path between LPARs. To simplify the illustration, only two logical CUs are shown.

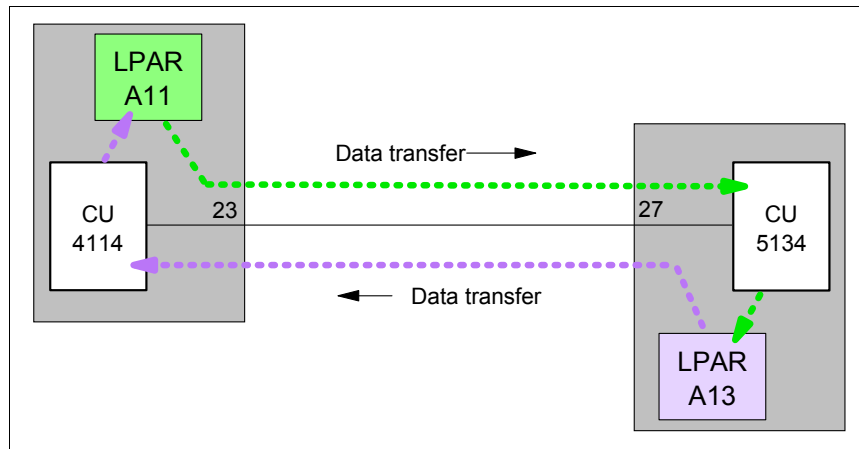


Figure 5-20 FCTC data transfer (point-to-point)

Data is transferred over the FICON link in both directions between the CUs and the logical partitions (LPARs):

- ▶ LPAR A11 sends data to LPAR A13 through CU 5134.
- ▶ LPAR A11 receives data from LPAR A13 through CU 4114.

The reverse applies to LPAR A13 when data is sent to or received from LPAR A11.

## 5.5 Building and activating the IODF

After the definition of the desired configuration is finished, a production IODF must be built and saved to the IODF volume and the IOCP data from the IODF written to an IOCDS on the Support Element. Up to four IOCDS files can be stored in a zEnterprise 196 server. This allows us to plan and define future I/O configurations and store them on the Support Element.

The tasks that are required to build a production IODF, save IOCP data, and activate changes dynamically are performed using HCD. Refer to the *HCD User's Guide*, SC33-7988, for detailed descriptions of all activation procedures.

HCM can be used to build the production IODF and save it to the IODF volume.

## 5.5.1 Building the production IODF

To build the production IODF:

1. Select the current work IODF using HCD. We are using SYS6.IODF05.WORK for our example.
2. From the HCD main menu, select **option 2.1 Build production I/O definition file**.
3. Review the Message List for any messages that require attention.
4. Press **PF3** to continue to the Build Production I/O Definition File panel.
5. Type the Production IODF name in the Production IODF name field. Refer Figure 5-21.

```
----- Build Production I/O Definition File -----

Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODF05.WORK'

Production IODF name . 'SYS6.IODF05' _____
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 5-21 HCD Build Production IODF

6. Press Enter and confirm the Descriptor field names. See Figure 5-22.

```
----- Define Descriptor Fields -----

Specify or revise the following values.

Production IODF name . . : 'SYS6.IODF05'

Descriptor field 1 . . . SYS6
Descriptor field 2 . . . IODF05
```

Figure 5-22 HCD Descriptor Fields

7. Press Enter to complete the build process. Message “Production IODF SYS6.IODF05 created.” confirms this.
8. Return to the HCD main menu, and select **option 2.11 Build and manage S/390® microprocessor IOCDs and IPL attributes**.
9. A list of processors that are accessible from the system that HCD is running on are displayed. Refer to Figure 5-23 on page 139. Type s next to the processor in our example.

```

S/390 Microprocessor Cluster List                               Row 1 of 6
Command ==>> _____ Scroll ==>> CSR

Select one or more CPCs, then press Enter.

-----CPC----- IODF
/ SNA Address      Type  Model  Processor ID
# IBM390PS.H40
# IBM390PS.TC8M
_ USIBMSC.SCZP101  2094  S18   SCZP101
_ USIBMSC.SCZP201  2097  E26   SCZP201
s USIBMSC.SCZP301  2817  M32   SCZP301
_ USIBMSC.SCZP901  2084  C24   SCZP901

```

Figure 5-23 HCD Writing an IOCDS from an IODF

The next panel shows a summary of all IOCDSs for this processor. Here we can see the current IOCDS that is in use by this processor, indicated by the YES-POR under the Write Protect column. Refer to Figure 5-24.

10.Type u next to IOCDS A2 to update this IOCDS with our new IODF information.

```

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.SCZP301  IODF77  LPAR  Alternate  No         No          No
_ A1.SCZP301  IODF00  LPAR  POR       Yes        No          Yes-POR
u A2.SCZP301  IODF75  LPAR  Alternate  No         No          No
_ A3.SCZP301  IODF76  LPAR  Alternate  No         No          No

```

Figure 5-24 HCD IOCDS List

11.Press enter to proceed to the next panel. Refer to Figure 5-25 on page 140.

12.Type IODF05 in the Title1 field. Leave the other options as defaults.

```

----- Build IOCDs -----
Row 1 of 1
Command ==> _____ Scroll ==> CSR

Specify or revise the following values.

IODF name . . . . . : 'SYS6.IODF05'

Title1 . IODF05
Title2 : SYS6.IODF05 - 2011-11-16 17:30

Write IOCDs in
IODCS      Switch IOCDs  preparation of upgrade
A2.SCZP301  No           No

```

Figure 5-25 HCD Build IOCDs

13. Press enter to continue to the final panel. See Figure 5-26.

```

----- Job Statement Information -----

Specify or revise the job statement information.

Job statement information
//IOCDSA2 JOB (ACCOUNT),'NAME',REGION=128M
//*
//*
//*
//*

```

Figure 5-26 HCD IOCDs Job Statement panel

14. Ensure that the JCL statements meet your site's requirements, and press enter to submit the Write IOCDs job.

15. After the job completes, verify the system output, and look for the following message. Refer to Example 5-5.

Example 5-5 HCD - IOCDs written successfully

---

```

CBDA674I IOCP successfully completed for A2.SCZP301.

```

---

16. Return to HCD Option 2.11, and navigate back to the IOCDs List for this processor. See Figure 5-27 on page 141.

Observe that IOCDs A2 now has a name of IODF05 and that there is now a Token Match of Yes under the IOCDs/Proc column.

```

                                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.SCZP301  IODF77  LPAR     Alternate  No         No           No
- A1.SCZP301  IODF00  LPAR     POR        Yes        No           Yes-POR
- A2.SCZP301  IODF05  LPAR     Alternate  No         Yes         No
- A3.SCZP301  IODF76  LPAR     Alternate  No         No           No

```

Figure 5-27 HCD IOCDS List after IODF update

## 5.5.2 Steps to activate the production IODF

There are two ways to activate an IODF. One way is to use the HCD panels in Option 2 and another way is using the ACTIVATE system command, which is what we demonstrate:

1. Access a console or SDSF under TSO for this system where commands are authorized to be entered.
2. Firstly enter **D IOS,CONFIG** to see what IODF the hardware and software are using for this system. Refer to Figure 5-28.

```

-D IOS,CONFIG
IOS506I 18.40.55 I/O CONFIG DATA 780
ACTIVE IODF DATA SET = SYS6.IODF00
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-10 09:25:30 SYS6 IODF00
ACTIVE CSS: 1 SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 5-28 MVS IODF system display

- Active IODF data set: Shows which IODF the Software (or LPAR) is currently using
  - Source: shows which IODF the Hardware (or HSA) for processor SCZP301 is currently using
3. Enter **ACTIVATE IODF=05, TEST** to test the new IODF, and verify that the changes we made to the IODF are reflected accurately in the test display. Refer Figure 5-29 on page 142. This command tests the changes that are being made to the Hardware or HSA.

```

-ACTIVATE IODF=05,TEST
IOS500I ACTIVATE RESULTS 675
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:
      1.4114-1.4117,1.5134-1.5137
      COMPID=SC1XL
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301:
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917
      COMPID=SC1XL
NOTE = A888,F
      COMPID=SC1XL
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 1.4114,1.5134
      COMPID=SC1XL
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301: 1.9000,1.9100
      COMPID=SC1XL

```

Figure 5-29 MVS Test ACTIVATE display

4. Enter **ACTIVATE, IODF=05, SOFT, TEST** to test changes being made to the software or operating system (TEST2094) for this system. Refer to Figure 5-30.

```

-ACTIVATE IODF=05,SOFT,TEST
IOS500I ACTIVATE RESULTS 718
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE
NOTE = 0100,SOFTWARE-ONLY CHANGE
      COMPID=SC1C3
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS WILL BE OUT OF SYNC
      COMPID=SC1C3

```

Figure 5-30 MVS Test ACTIVATE SOFT command

The responses to the displays look positive as per the message TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE.

We are now ready to activate the new IODF to the hardware and software for this system.

Because our system (SC30) is part of a sysplex (COMMPLEX), we must first activate the IODF on those systems. The systems defined in this sysplex are SC30, SC31, SC32, and SC33. Refer to Figure 5-31.

```

-D XCF
IXC334I 10.18.08 DISPLAY XCF 305
      SYSPLEX COMMPLEX:  SC30          SC31          SC32
                        SC33

```

Figure 5-31 MVS D XCF command

A SOFT ACTIVATE was performed already on systems SC32 and SC33.

5. Now for system SC31, enter **ACTIVATE IODF=05, SOFT**.

```

-RO SC31,ACTIVATE IODF=05,SOFT
IOS1500I ACTIVATE RESULTS 276
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 0100,SOFTWARE-ONLY CHANGE
  COMPID=SC1C3
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS ARE NOW OUT OF SYNC
  COMPID=SC1C3

```

Figure 5-32 MVS ACTIVATE SOFT command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 5-32.

- Now enter D IOS,CONFIG to verify that the activate for the LPAR SC31 completed.

```

-RO SC31,D IOS,CONFIG
IOS506I 10.19.07 I/O CONFIG DATA 283
ACTIVE IODF DATA SET = SYS6.IODF05
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-10 09:25:30 SYS6      IODF00
ACTIVE CSS: 1 SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 5-33 MVS D IOS.CONFIG command

Observe that the software is on SYS6.IODF05 and the hardware is on SYS6.IODF00. Refer to Figure 5-33.

- Now for system SC30, enter **ACTIVATE IODF=05**.

```

-RO SC30,ACTIVATE IODF=05
IOS1500I ACTIVATE RESULTS 364
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:
      1.4114-1.4117,1.5134-1.5137
  COMPID=SC1XL
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301:
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917
  COMPID=SC1XL
NOTE = A888,F
  COMPID=SC1XL
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 1.4114,1.5134
  COMPID=SC1XL
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301: 1.9000,1.9100
  COMPID=SC1XL

```

Figure 5-34 MVS ACTIVATE IODF command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 5-34.

- Enter D IOS,CONFIG to verify that the activate for the LPAR SC30 completed.

```

-RO SC30,D IOS,CONFIG
IOS506I 10.24.24 I/O CONFIG DATA 378
ACTIVE IODF DATA SET = SYS6.IODF05
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE    TIME    DESCRIPTION
SOURCE: SCZP301 11-11-16 16:39:06 SYS6    IODF05
ACTIVE CSS: 1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 5-35 MVS D IOS,CONFIG command

Observe that the software is on SYS6.IODF05 and the hardware is now on SYS6.IODF05. Refer to Figure 5-35.

The IOCDS must be switched to IOCDS A2 so that the Reset Profile for this processor SCZP301 points to the correct IODF information should there be a scheduled or unscheduled POR of this processor.

9. Enter **ACTIVATE ACTIOCDs=A2** to switch the IOCDS.

```

-RO SC30,ACTIVATE ACTIOCDs=A2
IOS500I ACTIVATE RESULTS 392
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 010C,THE ACTIVE IOCDS HAS BEEN SWITCHED TO A2
COMPID=SC1C3

```

Figure 5-36 MVS ACTIVATE ACTIOCDs command

Observe the message THE ACTIVE IOCDS HAS NOW BEEN SWITCHED TO A2. Refer to Figure 5-36.

10. If we now return to HCD, and select:

- SYS6.IODF05
- Option 2.11
- our test processor SCZP301

we see that IOCDS A2 is now selected and that a 'Yes' response is displayed for all columns. Refer to Figure 5-37.

```

                                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.SCZP301  IODF77    LPAR     Alternate   No         No           No
_ A1.SCZP301  IODF00    LPAR     Alternate   No         No           No
_ A2.SCZP301  IODF05    LPAR     POR         Yes        Yes          Yes-POR
_ A3.SCZP301  IODF76    LPAR     Alternate   No         No           No

```

Figure 5-37 HCD IOCDS List after new IODF activate

By pressing F20, more information is displayed in the IOCDS List panel. See Figure 5-38 on page 145.



```

IOCDs List                               Row 1 of 4 More: <
Command ==> _____ Scroll ==> CSR

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

--Last Update--
/ IOCDs      Date      Time  IOCDs Configuration Token Information
_ A0.SCZP301 2011-11-07 13:46 SCZP301 15:05:04 11-11-03 SYS6 IODF77
_ A1.SCZP301 2011-11-16 17:53 SCZP301 09:25:30 11-11-10 SYS6 IODF00
_ A2.SCZP301 2011-11-16 18:12 SCZP301 16:39:06 11-11-16 SYS6 IODF05
_ A3.SCZP301 2011-11-02 17:16 SCZP301 16:28:14 11-11-02 SYS6 IODF76

```

Figure 5-38 HCD IOCDs List after pressing F20

Now that the hardware component (HSA) of the IODF is synchronized, we must make sure that the software component is also synchronized.

The volume where the IODF was saved before (IODFPK) also contains a dataset called SYS#.IPLPARM. This dataset contains LOADxx members, which are used at IPL startup to determine certain system parameters. It also tells the system being IPLed which IODF to use.

When an LPAR is IPLed, a Load (IPL) address and a Load parameter (loadparm) are specified

Figure 5-39 shows what the Load Profile looks like for this system:

- ▶ The LOAD address points to where z/OS and programs reside
- ▶ The LOAD parameter C73001M1, where:
  - **C730** equals where the IODF resides and also the SYS#.IPLPARM dataset
  - **01** equals the suffix value of the LOADxx member in SYS#.IPLPARM
  - **M** equals the NIP suppression indicator
  - **1** equals alternate nucleus member

|   |  |  |
|---|--|--|
| Profile name                              | TEST9302   |  |
| Description                               | Z1DRB1 GA  |  |
| Load type                                 | <input type="radio"/> Normal <input checked="" type="radio"/> Clear <input type="radio"/> SCSI <input type="radio"/> SCSI_dump |  |
| Load address                              | 9302   | <input type="checkbox"/> Use dynamically changed address   |
| Load parameter                            | C73001M1   | <input type="checkbox"/> Use dynamically changed parameter |
| Time-out value                            | 300  | 60 to 600 seconds  |
| <input type="checkbox"/> Store status     |  |  |
| Worldwide port name                       | 0  |  |
| Logical unit number                       | 0  |  |
| Boot program selector                     | 0  |  |
| Boot record logical block address         | 0  |  |
| Operating system specific load parameters |  |  |

Figure 5-39 HMC Load Profile

11. Logon to the system where the IODF was written. In our example, it is SC30 and access SYS#.IPLPARM.

12. Enter the command **D IPLINFO** to obtain which LOADxx member the system used during IPL. Refer Figure 5-40 on page 146.

```

-D IPLINFO
IEE254I 14.34.15 IPLINFO DISPLAY 506
SYSTEM IPLED AT 22.03.35 ON 11/02/2011
RELEASE z/OS 01.13.00 LICENSE = z/OS
USED LOAD01 IN SYS0.IPLPARM ON OC730
ARCHLVL = 2 MTLSHARE = N
IEASYM LIST = XX
IEASYS LIST = (00) (0P)
IODF DEVICE: ORIGINAL(OC730) CURRENT(OC730)
IPL DEVICE: ORIGINAL(O9302) CURRENT(O9302) VOLUME(Z1DRB1)

```

Figure 5-40 MVS D IPLINFO command

Example 5-6 is a sample of the LOAD01 member from SYS0.IPLPARM, showing:

- ▶ HWNAME = SCZP301
- ▶ LPARNAME = A11
- ▶ IODF = 05

Example 5-6 HCD - LOAD01 member in SYS0.IPLPARM

---

```

*-----DEFINITION FOR SC30-----*
HWNAME SCZP301
LPARNAME A11
*IEASYM 99
SYSPLEX COMMPLEX Y
IODF 05 SYS6 TEST2094 01 Y
SYSCAT COMCAT123CMCAT.COMCAT
PARMLIB SYS1.PARMLIB
PARMLIB CPAC.ZOSR1D.PARMLIB
PARMLIB SYS1.IBM.PARMLIB
*-----DEFINITION FOR SC31-----*
HWNAME SCZP301
LPARNAME A13
*IEASYM 99
SYSPLEX COMMPLEX Y
IODF 05 SYS6 TEST2094 01 Y
SYSCAT COMCAT123CMCAT.COMCAT
PARMLIB SYS1.PARMLIB
PARMLIB CPAC.ZOSR1D.PARMLIB
PARMLIB SYS1.IBM.PARMLIB

```

---

This ensures that the operating system component of the IODF (software) when IPLed, matches the hardware component of the IODF (hardware) that is loaded into HSA.

## 5.6 Configuring the IBM Storage System DS8000

The host adapter ports in the DS8000 storage control unit must be configured for FICON (FC) protocol. Follow the procedure described in Appendix D, “Configuring the DS8000 for FICON” on page 415, to configure the host adapter ports.

For planning and implementation information, refer to *IBM System Storage DS8000: Architecture and Implementation*, SG24-8886.

## 5.6.1 Connecting the fiber optic cables

After all definition and configuration tasks are completed, fiber optic cables must be plugged into the FICON channels and DS8000 CU host bay adapters (HBAs).

All fiber links between the FICON channel and the CU host adapters use the longwave (LX) type of laser. For the LX type of laser, using a 9µm single mode (SM) fiber optic cable is recommended to achieve maximum performance and distance. All fiber optic cables must have LC duplex connectors at both ends.

Plug the fiber optic cable connectors to the designated ports, shown in Figure 5-1 on page 120, to the FICON channels and the CU ports, as listed here:

- ▶ PCHID 398 to DS8000 host adapter port 0203
- ▶ PCHID 3E0 to DS8000 host adapter port 0303

Make sure that the fiber optic cable ends are cleaned before plugging them into the transceivers. Also ensure that all fiber optic cables are labeled and documented for later use (for example, when recabling or for problem determination).

For further information and considerations regarding fiber cabling and documentation, refer to 4.11, “Physical connectivity” on page 113 and 4.2, “Documentation” on page 80.

## 5.6.2 Configuring the channel paths and devices online

After you successfully activate the new I/O definitions, you can configure the channel path online and try to communicate with the devices. Note that even if all the definitions in the IOCDs are correct, the channel status might be offline.

To configure the channel paths and devices online:

1. Enter **D M=CHP(20)** on the system console to display the current status of CHPID 20. If the CHPID is offline, enter **CF CHP(20),ON** to configure the channel online. The channel status must change to ONLINE, as shown in Example 5-7.

*Example 5-7 MVS - CF CHP(20) command*

---

```
-CF CHP(20),ONLINE
IEE502I CHP(20),ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

---

To achieve the best performance on the FICON channel, make sure that zHPF is enabled. Refer to 4.9.6, “High Performance FICON” on page 107 for considerations regarding how to exploit zHPF on System z servers.

2. Enter **D IOS,ZHPF** at the z/OS console to display zHPF settings for the z/OS image. If zHPF is disabled, enter **SETIOS ZHPF,YES**. This enables zHPF, as shown in Example 5-8.

*Example 5-8 MVS - Displaying zHPF*

---

```
-D IOS,ZHPF
IOS630I 15.14.09 ZHPF FACILITY 675
HIGH PERFORMANCE FICON FACILITY IS ENABLED
```

---

Using the **SETIOS ZHPF=YES** command enables zHPF temporarily. However, after the next system IPL, the zHPF facility, is reset to the default (disabled). To permanently enable zHPF for z/OS, add or change the ZHPF parameter in the SYS1.PARMLIB member IECIOSxx to ZHPF=YES. Refer to Example 5-9 on page 148.

Example 5-9 zHPF settings in SYS1.PARMLB(IECIOS00)

---

```
EDIT          SYS1.PARMLIB(IECIOS00) - 01.16          Columns 00001 00072
Command ==>                                         Scroll ==> CSR
***** ***** Top of Data *****
000001 HYPERPAV=YES
000002 ZHPF=YES
***** ***** Bottom of Data *****
```

---

3. Now you can query the status and functional details of the channel by entering **D M=CHP(20)** at the operator console. The command output is shown in Example 5-10. It provides information about the channel and the attached devices.

Example 5-10 MVS - D M=CHP(20)

---

```
-D M=CHP(20)
IEE174I 15.07.39 DISPLAY M 636
CHPID 20: TYPE=1A, DESC=FICON POINT TO POINT, ONLINE
DEVICE STATUS FOR CHANNEL PATH 20
      0 1 2 3 4 5 6 7 8 9 A B C D E F
0900 + + + + + + + + + + + + + + +
0901 + + + + + + + + + + + + + + +
0910 + + + + + + + + + + + + + + +
0911 + + + + + + + + + + + + + + +
...
1900 UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
1901 UL UL UL UL UL UL HA HA HA HA HA HA HA HA HA HA
1902 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1903 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1904 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
...
1917 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
SWITCH DEVICE NUMBER = NONE
ATTACHED ND = 002107.932.IBM.75.0000000L3001
PHYSICAL CHANNEL ID = 0398
FACILITIES SUPPORTED = ZHPF
```

---

Example 5-10 shows that CHPID 20 is online and operating in a point-to-point configuration. Information is also displayed about the attached devices and the facilities (for example: ZHPF) supported by the channel.

4. To verify that communication to the attached devices is working properly, enter **D M=DEV(xxxx)**, where xxxx is any device number. For example, to check the status of device 9000, which is a storage device, enter **D M=DEV(9000)** on a z/OS console. See Example 5-11 on page 149.

Example 5-11 MVS - D M=DEV(9000)

---

```
-D M=DEV(9000)
IEE174I 15.26.45 DISPLAY M 715
DEVICE 9000 STATUS=ONLINE
CHP          40   41   42   43   20   24
ENTRY LINK ADDRESS  22  22  1E  1E  ..  ..
DEST LINK ADDRESS  1B  1B  10  10  0D  0D
PATH ONLINE       Y   Y   Y   Y   Y   Y
CHP PHYSICALLY ONLINE Y   Y   Y   Y   Y   Y
PATH OPERATIONAL  Y   Y   Y   Y   Y   Y
MANAGED           N   N   N   N   N   N
CU NUMBER         9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED: 0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND         = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED     = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED    = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW, ZHPF
```

---

The response to **D M=DEV(9000)** displays all available paths to the devices and their status. Information about the device (for example, device type), the control unit number, and the functions supported by the device (for example, MIDAW and ZHPF) is shown.

### 5.6.3 Verifying the installation

After all the configuration tasks in the z196 server and the DS8000 storage are complete, you can verify that the current configuration matches the desired configuration shown in Figure 5-1 on page 120.

This verification must be done to ensure that the fiber optic cables are correctly plugged to the designated ports in the System z196 server, the FICON switches and the DS8000, and to check the status of channels and devices.

To verify the installation:

1. First, check the status of PCHID 398 in the z196 server. From any Hardware Management Console (HMC) where the desired CPC (server) is defined (for example, SCZP301), start a Single Object Operation (SOO) to that CPC's Support Element.

Display the list of channels installed on the CPC. Search the list for PCHID (398) and double-click the PCHID icon, which displays the PCHID details panel, as shown in Figure 5-41 on page 150.

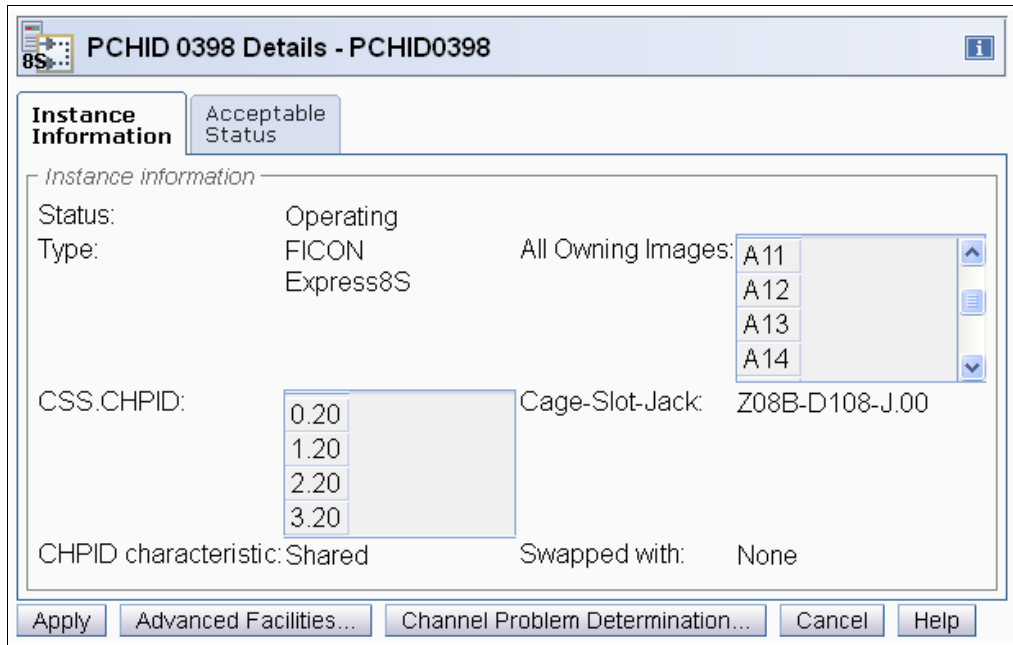


Figure 5-41 HMC PCHID details for 398

Essential information for PCHID 398 is shown on the PCHID Details panel:

- The PCHID status is Operating
- The PCHID type is FICON Express8S
- CSS.CHPID 1.20 is assigned to PCHID 398 (plus others)
- The owning images of PCHID 398 are A11 (plus others)
- The CHPID assigned to PCHID 398 is shared

2. The CHPID Details panel provides information that is similar to the PCHID details panel. To display the CHPID details, select the channel list for an LPAR from the Support Element. Refer to Figure 5-42.




Figure 5-42 HMC CHPID details for 1.20

Notice that the information for PCHID 398 (CHPID 1.20) provided on the details window matches the designated configuration in Figure 5-1 on page 120. This proves that PCHID 398 has CHPID 1.20 assigned and that image A11 can access the channel. Repeat these checks on other channels that were recently configured.

**Important:** If any of the data displayed on the PCHID or CHPID detail panel does not match the desired configuration, you must correct the definitions in the IOCDs.

3. Check that the FICON channels are connected to the correct host adapter port in the DS8000 storage controller. On either the PCHID details or CHPID details panel, click the **Channel Problem Determination** button. This displays the Channel Problem Determination panel, where you can select which information you want to display. See Figure 5-43.

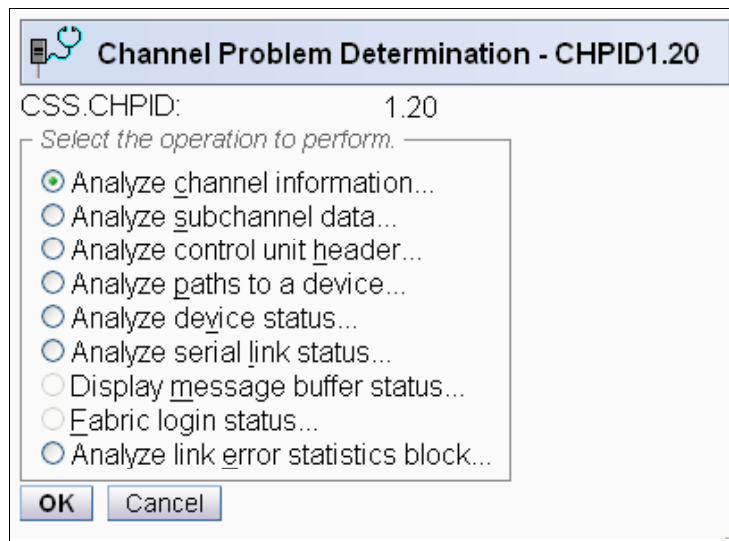


Figure 5-43 HMC Channel Problem Determination panel

4. Select **Analyze channel information**, and click **OK**. The Analyze Channel Information window is displayed, as shown in Figure 5-44 on page 152, which provides information about the node attached to the FICON channel.

| Analyze Channel Information - CHPID1.20  |                  |   |                        |
|--|------------------|---|------------------------|
| Channel type:  | FICON Native     | Hardware type:                            | E0                     |
| Partition ID:  | 11               | Hardware subtype:                         | 04                     |
| MIF image ID:  | 1                | 2 byte control unit link address defined: | No                     |
| Channel mode:  | Spanned          | Absolute address:                         | 000000009F4E8400       |
| CHPARM:  | 00               |   |                        |
| CSS.CHPID:   | 1.20             |   |                        |
| PCHID:   | 0398             |   |                        |
| Switch number:   | 00               |   |                        |
| Switch number valid:   | 0                | IFCC threshold:                           | 4                      |
| State:   | Online           | Channel link address:                     | 00                     |
| Status:  | Operating        | Temp error threshold:                     | 4                      |
| Image chnl state:  | Online           | Suppress:                                 | 0000000000000000       |
| Image chnl status:   | Operating        | SAP Affinity:                             | 05                     |
| Error code:  | A4               | Card description:                         | FICON Express8 10km LX |
| Ber inbound:   | 0                | Connection rate:                          | 4Gbps                  |
| Ber outbound:  | 0                |   |                        |
| Node type:   | Self             | Node type:                                | Attached               |
| Node status:   | Valid            | Node status:                              | Valid                  |
| Flag/parm:   | 10000720         | Flag/parm:                                | 00310100               |
| Type/model:  | 002817-M32       | Type/model:                               | 002107-932             |
| MFG:   | IBM              | MFG:                                      | IBM                    |
| Plant:   | 02               | Plant:                                    | 75                     |
| Seq. number:   | 0000000B3BD5     | Seq. number:                              | 0000000L3001           |
| Tag:   | F020             | Tag:                                      | 0203                   |
| World wide node name:  | 5005076400CB3BD5 | World wide node name:                     | 5005076305FFC74C       |
| World wide port name:  | 5005076028BF0F57 | World wide port name:                     | 500507630510C74C       |
| <input type="button" value="OK"/> <input type="button" value="Error Details..."/> <input type="button" value="Refresh"/> |                  |   |                        |

Figure 5-44 HMC Analyze Channel Information

Information about the nodes is displayed at the bottom part of the window. The lower left side displays information about the node in the z196 server. The lower right side displays information about the attached node.

**Node status:** Make sure that the Node status for both nodes is displayed as Valid. If any other status is shown, none of the displayed information is valid.

5. Check that the Type/model information and the serial number (Seq. number) is as expected.
6. Examine the Tag field for each node. The Tag provides information about the port number of the attached node. Under the z196 server, the two right-most digits of the Tag value represent the CHPID number for the channel node (20). For the attached node, the four digits represent the port number (0203). Be aware, however, that the tag value is provided by the attached device during link initialization and can have different meanings, depending on the vendor.



- The World Wide Node Name (WWNN) and the World Wide Port Name (WWPN) are also shown for each port, and can be used to prove that the channel is connected to the correct FICON adapter port if the WWNN or WWPN of the attached device is known.

If the node status is not Valid or the Tag value and WWPN value are not correct, check the fiber optic cable link between the z196 server and the DS8000 host adapter to ensure that it is plugged correctly.

In our scenario, we are now sure that PCHID 398 has connectivity to port number 0203 at the DS8000 host adapter, which matches our desired configuration. See Figure 5-1 on page 120.

If the displayed values are not as expected, the fiber optic cables might not be plugged correctly and must be checked.

After completing the preceding steps and proving that the physical path to the DS8000 storage controller and the logical definitions of the link are correct, check that the path to the control unit image is initialized correctly and properly defined.

- On the Channel Problem Determination panel, shown in Figure 5-43 on page 151, select **Analyze Serial Link Status**, and click **OK**. The Analyze Serial Link Status window is displayed, as shown in Figure 5-45.

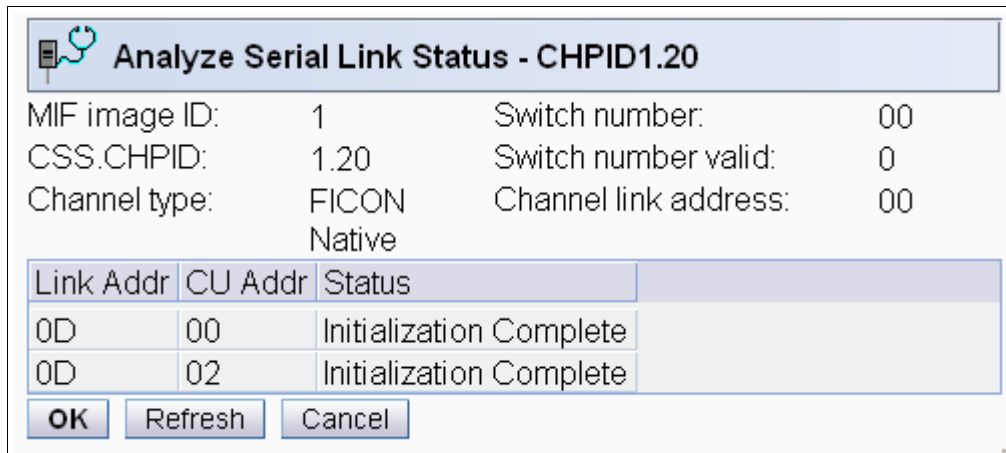


Figure 5-45 HMC Analyze Serial Link status

The Analyze Serial Link Status window provides status information about the link to the control unit images defined in the IODF. Scroll through the list of CU images, and check that the status for all CUs displays as **Initialization Complete**. Figure 5-45 shows a link status of **Initialization Complete** for all defined CU images on CHPID 20.

Although there is no FICON Director attached, a link address of **0D** is displayed by default for a point-to-point FICON link.

If the link status **Initialization Complete** is not shown, you must check that the ports in the CU are correctly configured and that the fiber optic cable link to the CU has the correct cable type and plugging.

Additionally the devices defined in the IODF might not have yet been defined in the DS8000.





## Configuring a switched topology

We describe all of the tasks that are required to configure and define a FICON environment for a switched topology.

Establishing a switched topology involves:

- ▶ Description of our environment
- ▶ Tasks and checklist:
  - Tasks
  - Verification checklist
- ▶ Defining the new configuration:
  - Defining the channel, CU, and storage devices
  - Defining our configuration example using HCD
  - Defining our configuration example using HCM
  - How the definition looks as IOCP statements
  - Defining FICON Channel-to-Channel (FCTC) connection
- ▶ Building and activating the IODF:
  - Building the production IODF
  - Activating the production IODF
- ▶ Configuring the IBM Storage System DS8000:
  - Configuring the FICON Director
  - Connecting the fiber optic cables
  - Configuring the channel path and device online
  - Verifying the installation

## 6.1 Establishing a switched topology

We describe the tasks that are required for configuring a FICON switched topology on a System zEnterprise 196 server, including setting up a FICON Director and configuring the attached storage controller. Our example is based on an infrastructure that is already built and installed, including the physical installation of:

- ▶ zEnterprise 196 server
- ▶ FICON Director
- ▶ DS8000 storage subsystem
- ▶ Fiber cabling infrastructure

## 6.2 Description of our environment

Based on considerations explained in Chapter 4, “Planning the FICON environment” on page 77, we are going to build a FICON switched configuration, which consists of an IBM zEnterprise 196 server, an IBM SAN768B-2 FICON Director, and an IBM DS8000 Storage subsystem.

In our example, all of the required infrastructure is already installed in our environment. The configuration is shown Figure 6-1.

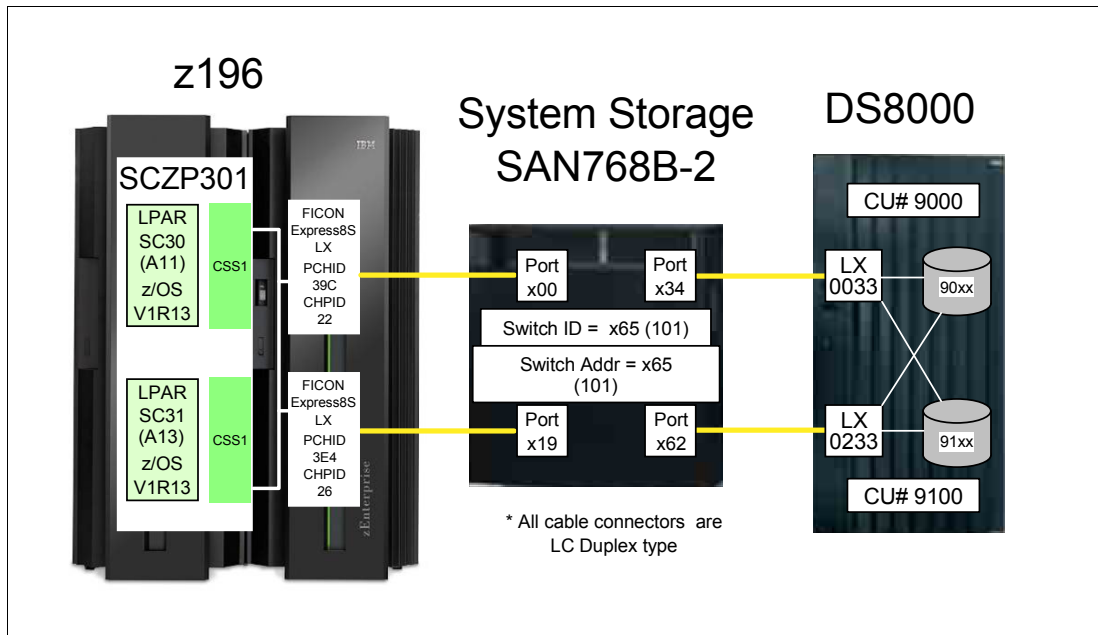


Figure 6-1 FICON Switched configuration

The z196 server (SCZP301) has two LPARs (A11 and A13) defined and activated. The system names running in these partitions are SC30 and SC31, respectively. The operating system running in both partitions is z/OS V1R13. As a standard feature of the z196 server, zHPF protocol is used for data transfer to the DS8000 storage device.

The two FICON channels are defined as spanned across CSS0, CSS1, CSS2, and CSS3 (only CSS1 is shown in our example). The channels are FICON Express8S LX features (FC 0409), using PCHID 39C CHPID 22 and PCHID 3E4 (CHPID 26).

The IBM SAN768B-2 FICON Director has connectivity to the z196 server on port x00 and port x19, and connectivity to the DS8000 storage system on port x34 and x62. All the ports in the FICON Director are longwave (LX) ports. The Switch ID (Switch #) and the Switch Address (Switch @) are both set to (hex) 65. The Switch Address is the Domain ID specified in the FICON Director and is defined as decimal, for example:

SwitchID / Switch Address of 65 (IODF) equals a Domain ID of 101 (switch)

The two host adapters installed in the DS8000 are longwave (LX) lasers. Port number 0033 and port number 0233 are used to connect to the FICON Director. Two logical control units (9000 and 9100) are defined and have devices 9000-90FF and 9100-91FF assigned. The DS8000 has the zHPF feature enabled.

A maximum unrepeated distance of 10 km (6.2 miles) is supported by the longwave laser (LX) feature when using 9µm single mode (SM). The fiber optic cables have an LC Duplex connector at both ends to connect to the z196 FICON Express8S channels: to the ports in the FICON Director and to the DS8000 host adapter.

**Note:** The SAN384B-2 and SAN768B-2 FICON Directors support 25km ISL links with FC2881 at 8Gbps LX.

## 6.3 Tasks and checklist

We provide an overview of the required tasks and a checklist that we used to make sure all hardware and software prerequisites were met before starting the configuration tasks.

### 6.3.1 Tasks

Figure 6-2 on page 158 shows the main steps required to define and activate a FICON switched configuration.

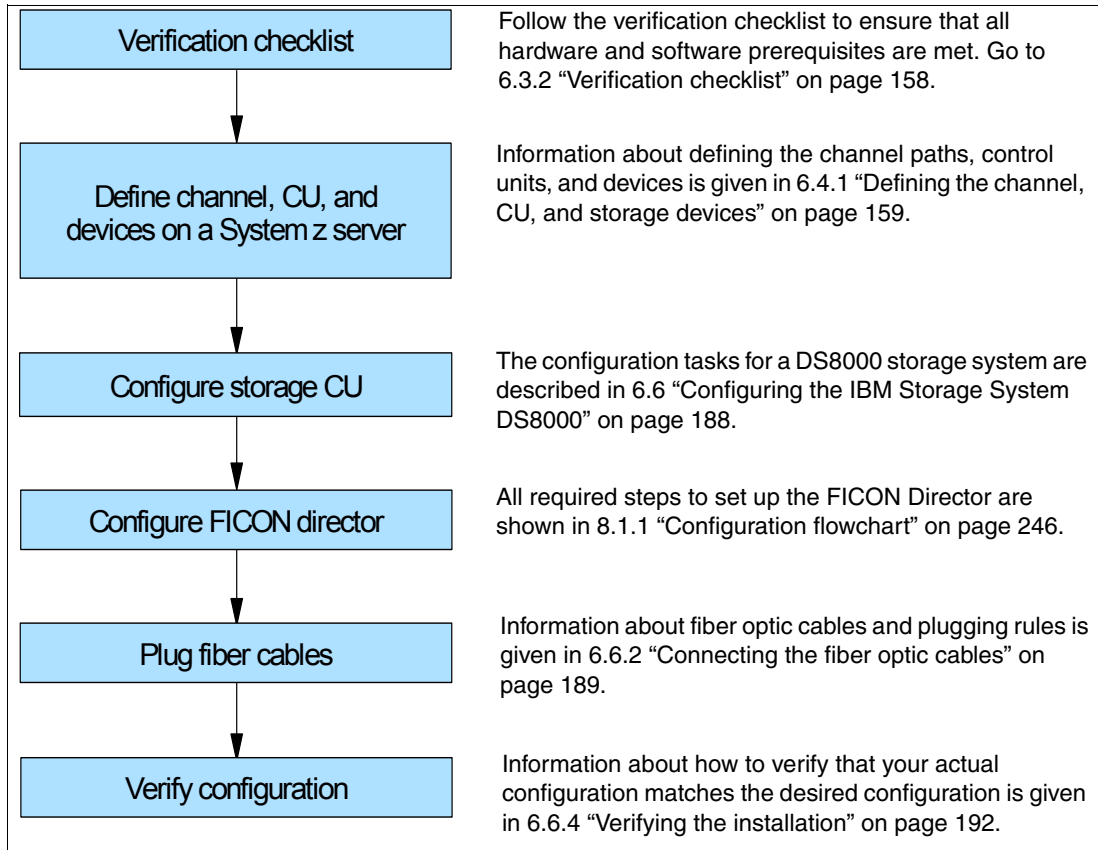


Figure 6-2 Steps for configuring and verifying a FICON switched configuration

### 6.3.2 Verification checklist

Before configuring the switched topology shown in Figure 6-1 on page 156, the following list was checked. All steps in the checklist must be finished and corrected (if required) to ensure a smooth and successful configuration of the topology. Hardware and software requirements must be checked.

- Check that the appropriate FICON features are available on the System z server. For details about each feature, refer to “System z FICON feature support” on page 37.
  - FICON Express4
    - LX FC 3321, FC 3324
    - SX FC 3322
  - FICON Express8
    - LX FC 3325
    - SX FC 3326
  - FICON Express8S
    - LX FC 0409
    - SX FC 0410
- If using the FICON Express8S feature, check the System z operating system requirements.
  - Check the 2817DEVICE PSP bucket for the latest information about FICON Express8S support on the operating system.

- ❑ Check that the number and types of FICON Director ports match the configuration requirements.
  - ❑ LX SFPs
  - ❑ SX SFPs
- ❑ Check DS8000 storage hardware requirements to support FICON longwave (LX) connectivity to the z10 server.
  - ❑ Check that FC 0709 and FC 7092 are installed to support zHPF.

**Note:** FC 0709 and FC 7092 must be ordered to obtain the license key to enable zHPF support in a DS8000 storage controller.

- ❑ Check that the DS8000 firmware is at level 6.1.16.0 or higher to support the zHPF feature.
- ❑ Check that the zHPF feature is enabled in the DS8000.
- ❑ Check that the correct fiber optic cables are available to connect the z196 server to the FICON Director, and the FICON Director to the DS8000 storage controller.
  - ❑ A 9µm single mode (SM) fiber optic cable is required to support longwave laser (LX) for a maximum distance of 10km (6.2 miles).
  - ❑ A 50µm or 62.5µm multi mode (MM) fiber optic cable is required to support shortwave laser (SX). See “System z FICON feature support” on page 37 for the maximum supported distance depending on cable type and speed.
  - ❑ An LC duplex connector is required at both ends of the fiber optic cables to connect to the z196 server FICON channel, the ports in the FICON Director, and the DS8000 host adapter ports.

**Note:** All fiber optic cables used in a link must be of the same type. For example, they must be either all single mode or all multi-mode fiber optic cables.

## 6.4 Defining the new configuration

This section describes all the tasks that are required to achieve the designated configuration shown in Figure 6-1 on page 156. It explains what needs to be done to get the FICON channels, FICON Director, storage control units, and the storage device online and operating.

For our scenario, we had an active partition in a zEnterprise 196 server running z/OS V1.13 with HCD. We used HCD to create, save, and activate our I/O definitions.

We also showed the definition process using HCM.

### 6.4.1 Defining the channel, CU, and storage devices

There are several tools available that you can use to define and configure a FICON environment on a System z server. The definition tools and an explanation of the statements and keywords used in the IOCP are described in Appendix C, “Configuration and definition tools” on page 409. Only statements and keywords relevant to our desired configuration are explained. Refer to *IOCP User's Guide*, SB10-7037, for details about all available IOCP statements and keywords.

Steps using both HCD and HCM are used to illustrate the definition process for the CHPIDs, CUs, and devices for the test configuration we describe in this section. Refer to *HCD User's Guide, SC33-7988*, or *HCM User's Guide, SC33-7989*, for more information and procedural example.

After all definition tasks are completed, the definition data is stored in the I/O definition file (IODF) data set. Any time you make changes to the I/O configuration in the "Work" IODF, a new "Production" IODF needs to be built and saved to the IODF volume (software) along with updating and writing to a spare IOCDS (hardware). This is required to maintain synchronization of the IODF between the software and hardware definitions.

The newly created "Production" IODF is used to write IOCDS information to the Support Element of the System z server. The Support Element in the System z server stores the IOCDS on its built-in HDD. The IOCDS stored in the Support Element is loaded into the server's storage (HSA) during a Power-on Reset (POR) of the System z server.

Changes or additions of channel paths, CUs, or devices can be performed dynamically in most cases. Because no system outage is required for dynamic changes, this is the preferred method for changing I/O definitions.

If changes to the I/O definitions are performed dynamically, those changes are directly updated into the Hardware System Area (HSA) of the server. The HSA is a separate storage area in the System z196 server that holds all I/O specifications and provides status information for all defined I/Os. The operating system uses the HSA to manage the I/O subsystem. For more details about System z196 channel subsystem characteristics, refer to 4.8.1 "System z" on page 97.

## 6.4.2 Defining our configuration example using HCD

To define our configuration example using HCD:

1. Select the current work IODF using HCD. We use SYS6.IODF06.WORK for our example.
2. From the HCD main menu, select **Option 1.3 Processor list**.
3. Type `s` next to the processor ID to define/verify the CHPID definitions (for example SCZP301).
4. Type `s` next to the CSS ID, which contains the LPARs (for example A11 and A13).
5. Type `c` next to our CHPIDs to 'Change Channel Path Definition' (for example CHPID 20). Refer to Figure 6-3 on page 161.



```

----- Change Channel Path Definition -----

Specify or revise the following values.

Processor ID . . . . : SCZP301
Configuration mode . : LPAR
Channel Subsystem ID : 1

Channel path ID . . . . 22   +           PCHID . . . 39C
Channel path type . . . FC   +
Operation mode . . . . . SPAN +
Managed . . . . . No (Yes or No)   I/O Cluster _____ +
Description . . . . . Express8s LX

Specify the following values only if connected to a switch:

Dynamic entry switch ID __ + (00 - FF)
Entry switch ID . . . . __ +
Entry port . . . . . __ +

```

Figure 6-3 HCD Change Channel Path Definition

In Figure 6-3, we can see that:

- Processor ID = SCZP301
- CSS = 1
- CHPID = 22
- PCHID = 39C

6. Press Enter, and observe the 'Define Access List' panel. Check that the required LPARs have access to CHPID 22. Refer to Figure 6-4.

```

----- Define Access List -----
Row 19 of 43
Command ==> _____ Scroll ==> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 1
Channel path ID . . . : 22   Channel path type . . : FC
Operation mode . . . : SPAN   Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
/ 1   A11                1    OS   COMPLEX SC30
/ 1    A12                 2     OS    VMLINUX9
/ 1   A13                3    OS   COMPLEX SC31
/ 1    A14                 4     OS
/ 1    A15                 5     OS
/ 1    A16                 6     OS    COMPLEX SC32
/ 1    A17                 7     OS    VMLINUX4
/ 1    A18                 8     OS    COMPLEX SC33
/ 1    A19                 9     OS    ITSOSS16
/ 2    A2A                 A     OS    ITSOSS13

```

Figure 6-4 HCD Define Access List

7. Perform the same steps for CHPID 26.
8. Return to the HCD main menu, and select **Option 1.2 Switches**.
9. Enter on the Command ==> **Add** to create a new switch definition.
10. Enter the highlighted values, as shown in Figure 6-5:
  - Switch ID = 65
  - Switch type = 2032
  - Switch address = 65
  - Installed port range = 00 - FF
  - Switch CU number = 0065
  - Switch device number = 0065

```

----- Add Switch -----

Specify or revise the following values.

Switch ID . . . . . 65 (00-FF)
Switch type . . . . . 2032      +
Serial number . . . . . _____
Description . . . . . SAN768B-2
Switch address . . . . . 65 (00-FF) for a FICON switch

Specify the port range to be installed only if a larger range
than the minimum is desired.

Installed port range . . 00 - FF +

Specify either numbers of existing control unit and device, or
numbers for new control unit and device to be added.

Switch CU number(s) . . . 0065  ____  ____  ____  ____  +
Switch device number(s) . 0065  ____  ____  ____  ____

```

Figure 6-5 HCD Add Switch

11. Press Enter to return to the Switch List.
12. Return to the HCD main menu, and select **Option 1.3 Processors**.
13. Navigate to CHPID 22 again in the Channel Path List.
14. Enter the highlighted values, as shown in Figure 6-6 on page 163:
  - Dyn SW = 65
  - Entry SW = 65
  - Entry Port = 00

```

Channel Path List      Row 25 of 123 More:  >
Command ==>> _____ Scroll ==>> CSR

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

Processor ID . . . . : SCZP301
Configuration mode . : LPAR
Channel Subsystem ID : 1

      PCHID          Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
- 22    39C   FC   SPAN  65 65 00      No Express8s LX
- 23    39D   FC   SPAN   _ _ _      No Express8s LX
- 24    3E0   FC   SPAN   _ _ _      No Express8s LX
- 25    3E1   FC   SPAN   _ _ _      No Express8s LX
- 26    3E4   FC   SPAN   _ _ _      No Express8s LX
- 27    3E5   FC   SPAN   _ _ _      No Express8s LX

```

Figure 6-6 HCD Define Switch Ports to Channel Path

15. Perform the same for our other test CHPID 26.
16. Return to the HCD main menu, and select **Option 1.4 Control units**.
17. Scroll down the list until CU 9000 is located.
18. Type c next to CU 9000 to “Change Control Unit Definition”. See Figure 6-7.
19. In our example, this Control Unit already contains existing definitions. We add connections to Switch #65 and CHPIDs 22 and 26 to these definitions.
20. Enter the highlighted values as shown in Figure 6-7:
  - Connected to switches = 65 65
  - Ports = 34 62

```

----- Change Control Unit Definition -----

Specify or revise the following values.

Control unit number . . . . 9000 +
Control unit type . . . . . 2107      +
Serial number . . . . . 0L30002421 +
Description . . . . . YU CU 0

Connected to switches . . . 61 62 63 64 65 65 _ _ +
Ports . . . . . 1B 1B 10 10 34 62 _ _ +

Define more than eight ports . . 2    1. Yes
                                   2. No

```

Figure 6-7 HCD Change Control Unit Definition

21. Press Enter to go to the Processor / CU panel.

22. Enter the highlighted values as shown in Figure 6-8:

- Channel Path ID. Link Address = 22.34 and 26.62

```

                                Select Processor / CU   Row 7 of 23 More:   >
Command ==> _____ Scroll ==> CSR

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

Control unit number . . . : 9000      Control unit type . . . : 2107

-----Channel Path ID . Link Address + -----
/ Proc.CSSID 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
_ SCZP301.0  40.1B  41.1B  42.10  43.10  _____  _____  _____  _____
_ SCZP301.1  40.1B  41.1B  42.10  43.10  20         24         22.34  26.62
_ SCZP301.2  40.1B  41.1B  42.10  43.10  _____  _____  _____  _____
_ SCZP301.3  40.1B  41.1B  42.10  43.10  _____  _____  _____  _____

```

Figure 6-8 HCD Adding CHPIDs 22 and 26 to SCZP301.1, CU 9000

23. Press Enter to observe the Modify Device Parameters panel. See Figure 6-9.

```

                                Select Processor / CU
----- Modify Device Parameters -----
                                Row 1 of 3 More:   >
Command ==> _____ Scroll ==> CSR

Specify or revise any changes to the device parameters in the list below.
To view attached control units, scroll to the right.

Processor ID . . . . : SCZP301
Channel Subsystem ID : 1

-----Device-----      ---UA----      Preferred Exposure
No., Range Type          SS+ Old New + Time-Out STADET CHPID + Device
9000,016  3390B          _  00 00  No    Yes    _
9000,128  3390A          1  80 80  No    Yes    _
9010,112  3390B          _  10 10  No    Yes    _
***** Bottom of data *****

```

Figure 6-9 HCD Modify Device Parameters

This panel will display the Unit Addresses (UA) already defined for the device:

- Devices 9000-907F as 3390B (SubChannelSet or SS=0)
- Devices 9000-907F as 3390A (SubChannelSet or SS=1)

**Note:** With the addition of HyperPAV, we do not need to define 128 alias devices. Refer 4.9.4 “Parallel Access Volume and HyperPAV” on page 105 to calculate your alias requirements.

24. Return to the HCD main menu, and select **Option 1.1 Operating system configurations.**

25. Type u next to the Operating system for our example (TEST2094)

26. Scroll down the list until I/O Device 9000 is located.

27. Type **c** next to 9000, and press Enter until the Define Device to Operating System Configuration panel is displayed.
28. Type **s** next to TEST2094 to start the Device Parameters / Features panel. See Figure 6-10.

```

----- Define Device to Operating System Configuration -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Select OSs to connect or disconnect devices, then press Enter.

Device number . : 9000           Number of devices : 16
Device type   . : 3390B

/ Config. ID  Type   SS Description          Defined
_ ALLDEV      MVS    All devices             Yes
_ LABSERV1    MVS    Lab Services           Yes
_ LO6RMVS1    MVS    Sysplex systems       Yes
_ MVSW1       MVS    Production systems     Yes
_ OPENMVS1    MVS    OpenEdition MVS       Yes
s TEST2094   MVS    Sysplex systems       Yes

```

Figure 6-10 HCD Defining a Device to an Operating System Configuration

Here we can enable/disable the device specific features by over typing the yes/no values. Refer Figure 6-11.

```

----- Define Device Parameters / Features -----
                                                    Row 1 of 6
Command ==> _____ Scroll ==> CSR

Specify or revise the values below.

Configuration ID . : TEST2094   Sysplex systems
Device number   . . : 9000      Number of devices : 16
Device type     . . . : 3390B

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 *****

```

Figure 6-11 HCD Define Device Parameters / Features

29. Press Enter to go to the Assign/Unassign Device to Esoteric panel. See Figure 6-12 on page 166.

```

----- Assign/Unassign Device to Esoteric -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Specify Yes to assign or No to unassign.  To view devices already
assigned to esoteric, select and press Enter.

Configuration ID : TEST2094           Sysplex systems
Device number   . : 9000             Number of devices  : 16
Device type    . . : 3390B          Generic . . . . . : 3390

+-oteric Assigned Starting Number Number of Devices
- 01.ATL3      No      _____ _____
- 01.CART      No      _____ _____
- 01.NOVIO     Yes    _____ _____
- 01.SYSDA     Yes    _____ _____
- 01.SYSSQ     No      _____ _____
- 01.TAPE      No      _____ _____
- 01.TAPE3592 No      _____ _____
- 01.VIO       Yes    _____ _____

```

Figure 6-12 HCD Assign/Unassign Device to Esoteric

This panel allows us to assign devices to particular esoterics, such as SYSDA and VIO.

30. Perform the same steps for the remainder of the devices for this control unit (3390B 9010-7F and 3390A 9000-7F), and then return to the HCD main menu.

Steps 16 through 30 can now be performed for Control Unit 9100 and Devices 9100-917F.

### 6.4.3 Defining our configuration example using HCM

To define our configuration example using HCM:

1. Select the current work IODF using HCM. We use SYS6.IODF06.WORK for our example. See Figure 6-13 on page 167.




Figure 6-13 HCM File Open IODF06.hcm

2. Select **Locate** → **Processor**.
3. Double-click the required processor in the list (for example SCZP301). See Figure 6-14.




Figure 6-14 HCM Processor display

Our configuration already contains existing definitions, so we will be adding to these.

4. Double-click the top pane of the processor display to edit the processor, or select **Edit** → **Processor**, and select the processor from the list. Refer to Figure 6-15 on page 168.

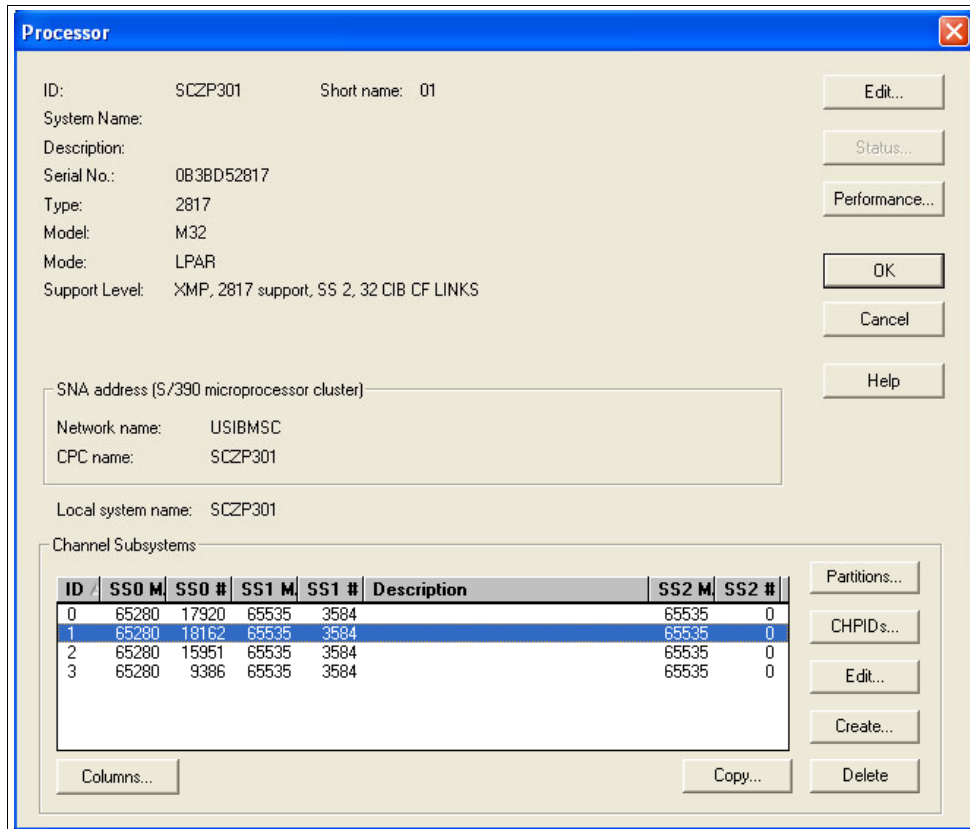


Figure 6-15 HCM Edit Processor

5. Highlight CSS ID 1, and click **CHPIDs**.
6. Scroll down the list until CHPIDs 22 and CHPIDs 26 are located. See Figure 6-16.

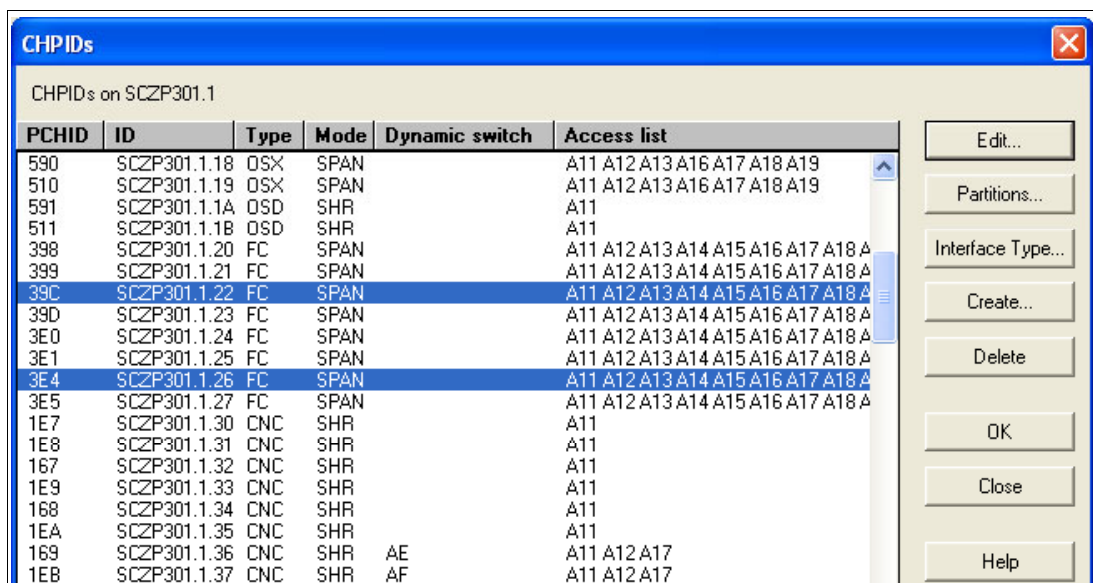
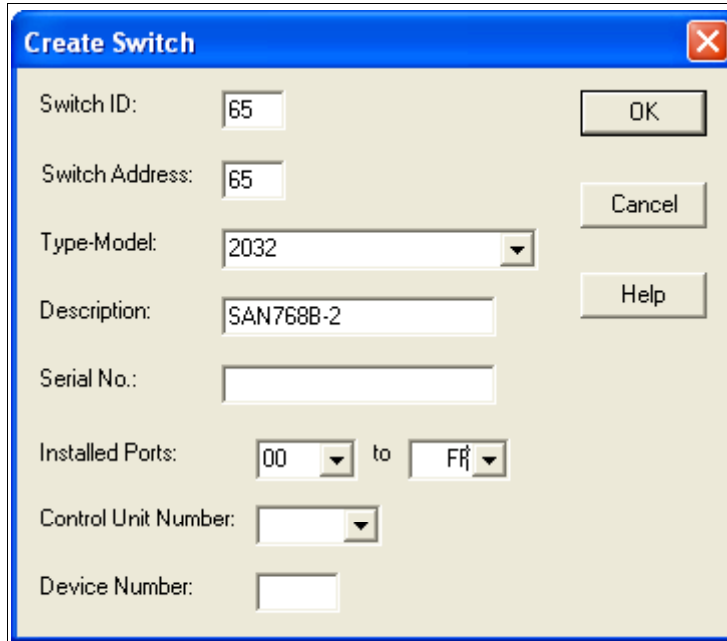


Figure 6-16 HCM Edit CHPIDs

7. Verify that they are defined correctly and have access to our test LPARs A11 and A13.



8. Select **Create** → **Switch**, and enter the following values, as shown in Figure 6-17:
- Switch ID = 65
  - Switch Address = 65
  - Type-Model = 2032
  - Installed Ports = 00 to FF
  - Control Unit Number = 0065
  - Device Number = 0065



The screenshot shows a 'Create Switch' dialog box with the following fields and values:

- Switch ID: 65
- Switch Address: 65
- Type-Model: 2032
- Description: SAN768B-2
- Serial No.: (empty)
- Installed Ports: 00 to FF
- Control Unit Number: (empty)
- Device Number: (empty)

Buttons: OK, Cancel, Help

Figure 6-17 HCM Create Switch

9. Click **OK**.
10. Double-click the switch object to list all the defined ports.
11. Highlight Port 00. Click **Connect**.
12. Navigate through the processors and CHPIDs until you find CHPID 22.
13. Highlight CHPID 22. Click **Connect** → **Physical Only**.
14. Perform the same action for CHPID 26 and Port 19. See Figure 6-18 on page 170.




Figure 6-18 HCM Ports connected to Channel Paths

15. Click **OK**, select **Edit** → **Controller**, and select the DASD controller for our test. See Figure 6-19 on page 171.

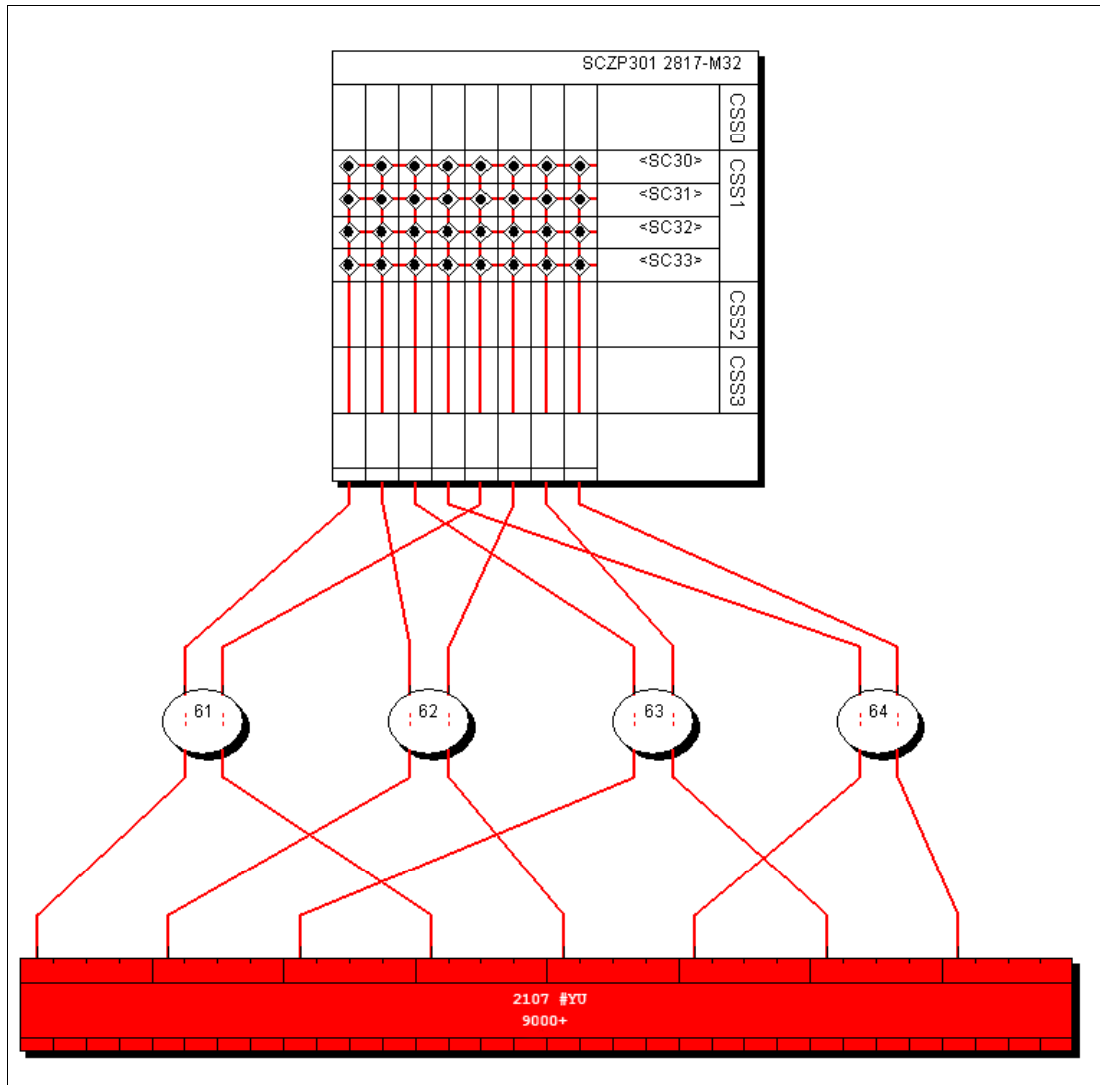


Figure 6-19 HCM DASD Controller view

Presently our DASD subsystem has eight FICON connections spread over four FICON directors plus two point-to-point connections. For our example we add two new switched FICON connections through Switch ID #65 to the DS8000.

16. Double-click the controller, and in the Interfaces box highlight our two interfaces 0033 and 0233. Click **Connect**.
17. Highlight the first controller interface in the left panel (0033).
18. Locate Switch #65 in the right panel, and expand the port list until port 34 is located.
19. Highlight port 34, and click **Connect**.
20. Perform the same action for Controller interface 0233. See Figure 6-20 on page 172.




Figure 6-20 HCM Connecting to DS8000 interfaces to Switch Ports

21. Click **Standard** to continue to the Affected Control Units panel.
22. Highlight CU 9000, and then click **Connect**.
23. Highlight CHPID 22, and then click **Add**. See Figure 6-21.

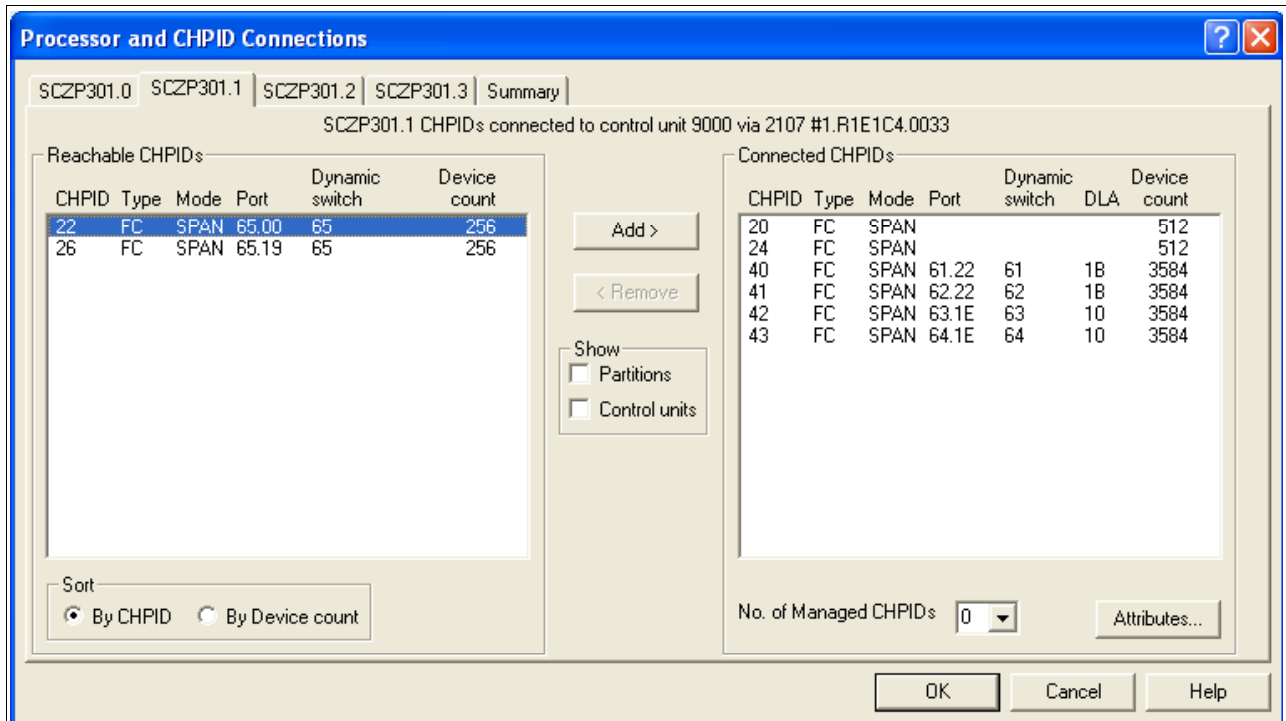


Figure 6-21 HCM Adding CHPIDs to CU definitions

24. Click **OK** to perform the same action for CU 9100.

25. Click **Close** to add CHPID 26 to CU 9100 and CU9100 through interface 0233. The final step looks like Figure 6-22.

Observe the 2 byte Destination Link Addresses (DLA). Refer Figure 6-22.

The connections marked in red show:

- CHPID 22 -> Switch #65 / Port 00 -> Switch #65 / Port 34 -> CU interface 0233
- CHPID 26 -> Switch #65 / Port 19 -> Switch #65 / Port 62 -> CU interface 0233

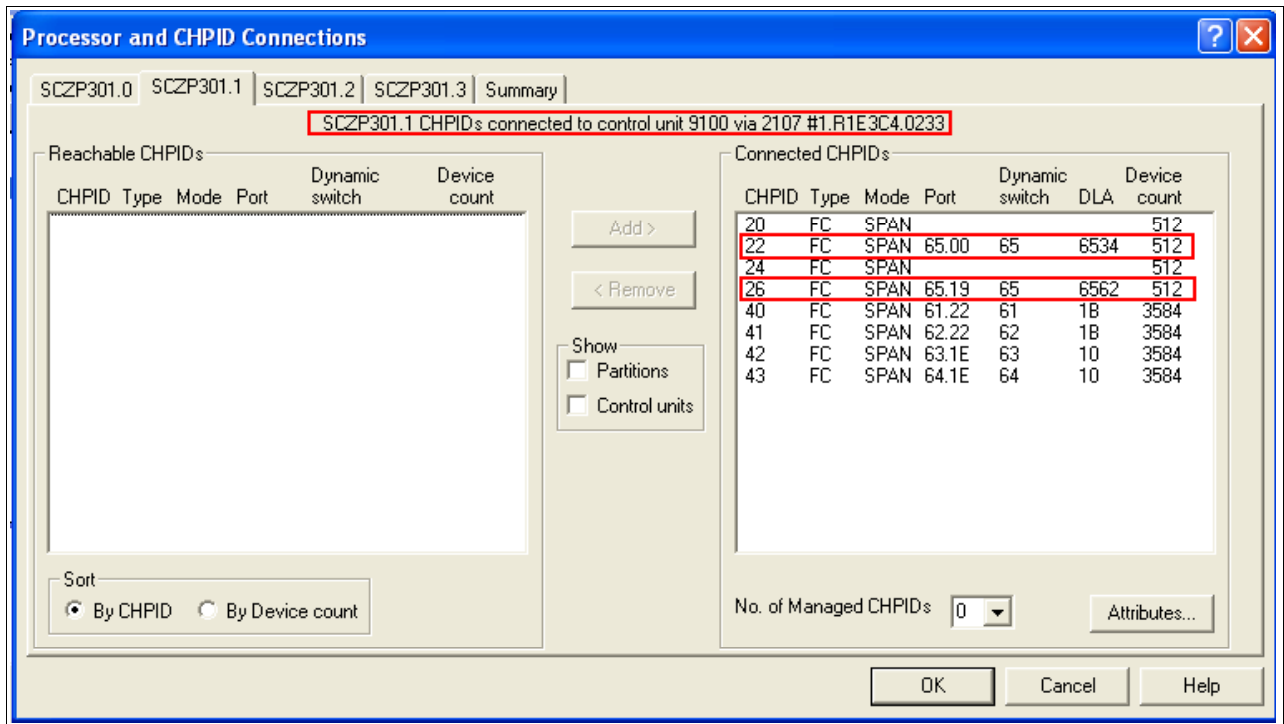


Figure 6-22 HCM Final step with adding CHPIDs to CU definitions

26. Click **OK** to return to the main HCM display and observe the new connections into interfaces 0033 and 0233. Refer Figure 6-23.




Figure 6-23 HCM DASD switched connections

Now we must ensure that the device definitions for the Operating System (software) are correct.

27. Double-click the controller box.

28. Click **Control Units**.

29. Highlight CU 9000, and then click **I/O Devices**.

30. Highlight all of the 3390B devices (9000-907F), and then click **OS Configs**.

31. Highlight TEST2094, and then click **Define**. Refer to Figure 6-24.

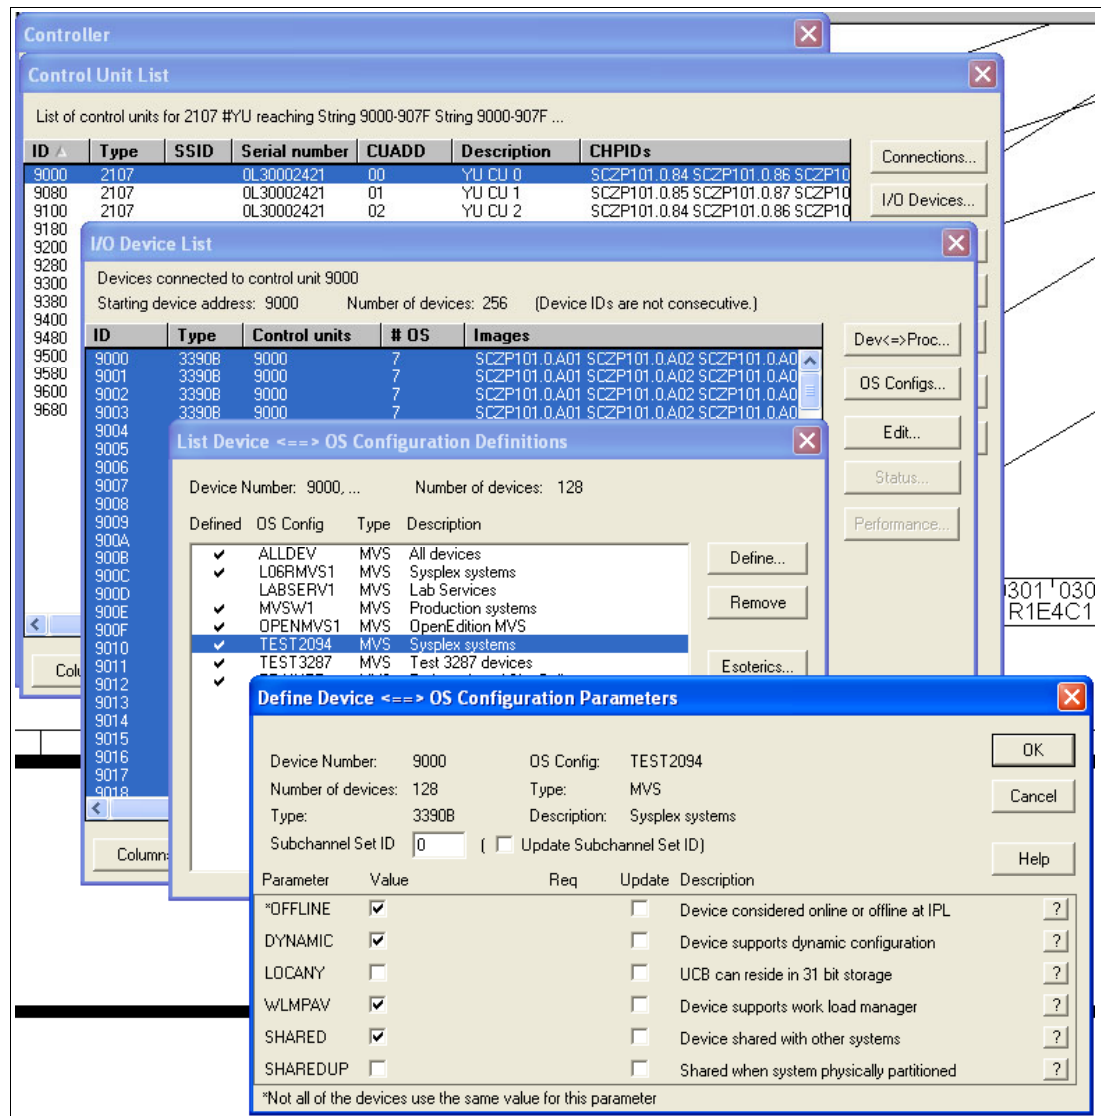


Figure 6-24 HCM OS Configuration Parameters

32. Verify that the Device Parameter values are correct.

33. Click **OK**. Click **Esoterics**, and verify that the devices are defined to the correct Esoterics.

34. Perform the same operation for the 3390A devices, CU 9000, CU 9100, and devices 9100-917F.

## 6.4.4 How the definition looks as IOCP statements

The following examples show the CHPID, CNTLUNIT, and IODEVICE statements when all the definition tasks based on the configuration shown in Figure 6-1 on page 156 are finished.

First we show the CHPID statement and its keywords. See Example 6-1

*Example 6-1 CHPID definition for switched configuration*

---

```

CHPID PATH=(CSS(0,1,2,3),22),SHARED, *
      PARTITION=((CSS(3),(A31,A34,A35),(=))), *
      NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),(=)),(CSS(1), *
      (A1B,A1E,A1F),(=)),(CSS(2),(A2E,A2F),(=))),SWITCH=65, *
      PCHID=39C,TYPE=FC
CHPID PATH=(CSS(0,1,2,3),26),SHARED, *
      PARTITION=((CSS(3),(A31,A34,A35),(=))), *
      NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),(=)),(CSS(1), *
      (A1B,A1E,A1F),(=)),(CSS(2),(A2E,A2F),(=))),SWITCH=65, *
      PCHID=3E4,TYPE=FC

```

---

- ▶ The PATH keyword in the CHPID statement defines CHPID 22 and CHPID 26 in CSS(1) as SHARED.
- ▶ The NOTPART keyword specifies the LPARs which do not have access to CHPID 22 and CHPID 26
- ▶ The SWITCH keyword shows that this CHPID is connected to Switch #65
- ▶ The PCHID keyword specifies the physical channel identification number associated with CHPID 22 and CHPID 26

With the CHPIDs defined, we next show how the Control Units attached to the CHPIDs are defined; Example 6-2 displays the CNTLUNIT statement and keywords.

*Example 6-2 CNTLUNIT definition for point-to-point configuration*

---

```

CNTLUNIT CUNUMBR=0065,PATH=(CSS(1),22,26), *
      UNITADD=((00,001)),LINK=(CSS(1),65FE,65FE),UNIT=2032
CNTLUNIT CUNUMBR=9000, *
      PATH=(CSS(0),40,41,42,43),(CSS(1),40,41,42,43,20,24,22, *
      26),(CSS(2),40,41,42,43),(CSS(3),40,41,42,43)), *
      UNITADD=((00,256)), *
      LINK=(CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,**,**,653* *
      4,6562),(CSS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)), *
      CUADD=0,UNIT=2107
CNTLUNIT CUNUMBR=9100, *
      PATH=(CSS(0),40,41,42,43),(CSS(1),40,41,42,43,20,24,22, *
      26),(CSS(2),40,41,42,43),(CSS(3),40,41,42,43)), *
      UNITADD=((00,256)), *
      LINK=(CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,**,**,653* *
      4,6562),(CSS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)), *
      CUADD=2,UNIT=2107

```

---

- ▶ The CUNUMBR keyword shows definitions for 0065, 9000 and 9100.
- ▶ The PATH keyword shows CSS(1) as having definitions for CHPID 22 and CHPID 26
- ▶ The LINK keyword shows the 2 byte switch and destination link address (DLA) that the control units have been defined to
- ▶ The CUADD keyword (representing the logical control unit or LCU) is set to 0 for CU 9000 and 2 for 9100. By pressing F20 in the panel showed in Figure 6-8 on page 164, you can see the CUADD value set.

After the CHPIDs and the CUs are defined, the next step is to define the devices owned by the CUs; Example 6-3 displays the IODEVICE statement and keywords.

*Example 6-3 IODEVICE definition for point-to-point configuration*

---

```
IODEVICE ADDRESS=065,UNITADD=00,CUNUMBR=(0065),STADET=Y,      *
          UNIT=2032
IODEVICE ADDRESS=(9000,128),CUNUMBR=(9000),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9000,128),UNITADD=80,CUNUMBR=(9000),      *
          STADET=Y,SCHSET=1,UNIT=3390A
IODEVICE ADDRESS=(9100,128),CUNUMBR=(9100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9100,128),UNITADD=80,CUNUMBR=(9100),      *
          STADET=Y,SCHSET=1,UNIT=3390A
```

---

- ▶ The ADDRESS keyword shows device address 065 for the Switch CUP device along with the DS8000 devices 9000 for 128 (9000-907F) or 9100 for 128 (9100-917F)
- ▶ The CUNUMBR keyword associates the IODEVICE statement with the Control Unit number 9000 or 9100
- ▶ The UNIT keyword shows the device type of 3390B (for the base devices) or 3390A (for the HyperPAV alias devices)
- ▶ The UNITADD keyword is used when the Unit Address (UA) does not start at 00
- ▶ The SCHSET keyword shows that these devices are defined in Subchannel Set 1 (SS1)

**Note:** For CU 9000 and CU 9100, the device allocation is defined in the following way

- ▶ 3390B devices 9000-907F, CUADD=0, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9000-907F, CUADD=0, UA = 80-FF, SCHSET = 1
- ▶ 3390B devices 9100-917F, CUADD=2, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9100-917F, CUADD=2, UA = 80-FF, SCHSET = 1

## 6.4.5 Defining FICON Channel-to-Channel (FCTC) connection

In addition to storage control units and storage devices and the CUPs in the FICON Directors, we will be defining some switched FICON Channel-to-Channel (FCTC) CUs and devices to allow communication between LPARs on the same processor. For more information about considerations regarding FCTCs, refer to Appendix G, “Adding FICON CTC connections” on page 489.

Based on those considerations, we are going to configure FCTCs in a switched environment. Figure 6-25 on page 177 illustrates our configuration. To simplify the illustration only four CUs and two LPARs are shown in this sample configuration, although more than four CUs would normally be defined.






Figure 6-25 FCTC configuration (switched)

In this configuration, there are two LPARs (A11 and A13) residing in the same z196 server. Both LPARs can access CHPIDs 22 and 26 because of shared definition of those channel paths.

Channels are attached to FICON Director 65 on ports 00 and 19. For FCTC communication, there are four logical control units defined:

- ▶ Two of them for each partition for local communication in the same LPAR
- ▶ Two of them for communication to the other LPAR over the FICON Director link

Each LPAR has access to a logical control unit and their logical devices defined in the corresponding LPAR:

- ▶ LPAR A11 will communicate through a FICON Director channel to CU 5134, which is defined to LPAR A13.
- ▶ LPAR A11 will communicate through a FICON Director channel to CU 5114, which is defined to LPAR A11.
- ▶ LPAR A13 will communicate through a FICON Director channel to CU 4114, which is defined to LPAR A11.
- ▶ LPAR A13 will communicate through a FICON Director channel to CU 4134, which is defined to LPAR A13.

Example 6-4 on page 178 shows the definition of FCTC control units and device definitions:

- ▶ CUNUMBR keyword shows CU 411C and 511C for LPAR A11 and CU 413C and 513C for LPAR A13.
- ▶ PATH keyword shows the CSS and CHPID that the CU is defined to.

- ▶ CUADD keyword shows the logical control unit number (LCU), which must equal the LPAR ID.
- ▶ NOTPART keyword shows the partitions that cannot access these devices. This is controlled using the Explicit Device Candidate List (EDCL) option in the operating system configuration definition.

*Example 6-4 IOCDs definitions for FICON CTC*

---

```

CNTLUNIT CUNUMBR=411C,PATH=((CSS(1),22)),UNITADD=((00,004)), *
LINK=((CSS(1),6519)),CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(411C,004),UNITADD=00,CUNUMBR=(411C), *
STADET=Y,NOTPART=((CSS(1),A11,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=413C,PATH=((CSS(1),22)),UNITADD=((00,004)), *
LINK=((CSS(1),6519)),CUADD=13,UNIT=FCTC
IODEVICE ADDRESS=(413C,004),UNITADD=00,CUNUMBR=(413C), *
STADET=Y,NOTPART=((CSS(1),A13,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=511C,PATH=((CSS(1),26)),UNITADD=((00,004)), *
LINK=((CSS(1),6500)),CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(511C,004),UNITADD=00,CUNUMBR=(511C), *
STADET=Y,NOTPART=((CSS(1),A11,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=513C,PATH=((CSS(1),26)),UNITADD=((00,004)), *
LINK=((CSS(1),6500)),CUADD=13,UNIT=FCTC
IODEVICE ADDRESS=(513C,004),UNITADD=00,CUNUMBR=(513C), *
STADET=Y,NOTPART=((CSS(1),A13,A1B,A1E,A1F)),UNIT=FCTC

```

---

There are four devices defined for each LPAR (LPAR A11 and LPAR A13), which allows communications between both LPARs using their corresponding devices:

- ▶ LPAR SC30 (A11) communicates with LPAR SC31 (A13) using devices 513C-513F
  - ▶ LPAR SC30 (A11) communicates with LPAR SC30 (A11) using devices 511C-511F
  - ▶ LPAR SC31 (A13) communicates with LPAR SC30 (A11) using devices 511C-511F
  - ▶ LPAR SC31 (A13) communicates with LPAR SC31 (A13) using devices 513C-513F
- Devices 411C-F and 511C-F cannot be brought online to LPAR SC30 (A11)
  - Devices 413C-F and 513C-F cannot be brought online to LPAR SC31 (A13)

The rules we followed for device numbering are described in “FCTC device numbering scheme” on page 490.

Figure 6-26 on page 179 illustrates the logical view of the FCTC configuration and the data path between two corresponding FCTC devices. To simplify the illustration, only two logical CUs are shown.

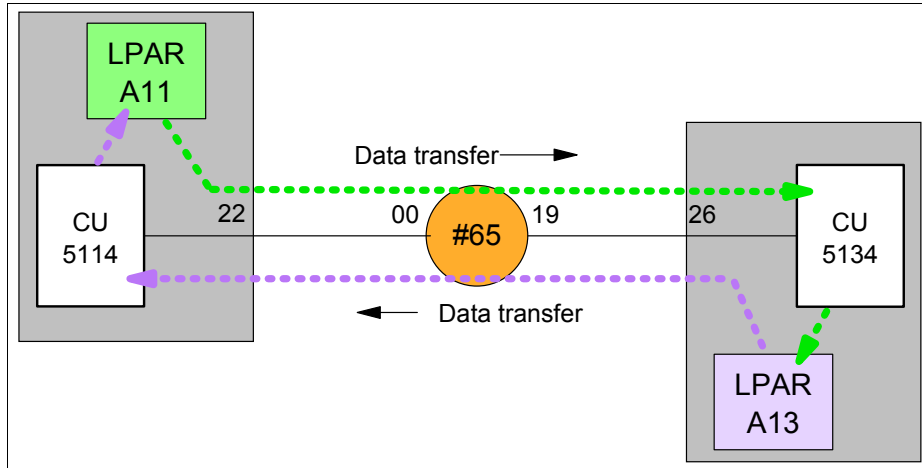


Figure 6-26 FCTC Data transfer (switched)

Data is transferred over the FICON link in both directions between the CUs and the logical partitions (LPARs):

- ▶ LPAR A11 sends data to LPAR A13 through CU 5134.
- ▶ LPAR A11 receives data from LPAR A13 through CU 5114.

The reverse applies to LPAR A13 when data is sent to or received from LPAR A11.

## 6.5 Building and activating the IODF

After the definition of the desired configuration is finished, a production IODF needs to build and save to the IODF volume, and the IODF data must be set and written to an IOCDS on the Support Element. Up to four IOCDS files can be stored in a zEnterprise 196 server. This allows us to plan and define future I/O configurations and store them on the Support Element.

The tasks that are required to build a production IODF, save I/O definition data, and activate changes dynamically are performed through HCD. Refer to the *HCD User's Guide*, SC33-7988, for detailed descriptions of all activation procedures.

### 6.5.1 Building the production IODF

To build the production IODF:

1. Select the current work IODF using HCD. We use SYS6.IODF06.WORK for our example.
2. From the HCD main menu, select **Option 2.1 Build production I/O definition file**.
3. Review the Message List for any messages that require attention.
4. Press PF3 to continue to the Build Production I/O Definition File panel.
5. Type the Production IODF name in the Production IODF name field. Refer to Figure 6-27 on page 180.

```

----- Build Production I/O Definition File -----

Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODF06.WORK'

Production IODF name . 'SYS6.IODF06' _____
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 6-27 HCD Build Production IODF

6. Press enter, and confirm the Descriptor field names. See Figure 6-28.

```

----- Define Descriptor Fields -----

Specify or revise the following values.

Production IODF name . : 'SYS6.IODF06'

Descriptor field 1 . . . SYS6
Descriptor field 2 . . . IODF06

```

Figure 6-28 HCD Descriptor Fields

7. Press Enter to complete the build process. The message Production IODF SYS6.IODF06 created confirms completion of the build process.
8. Return to the HCD main menu and select **Option 2.11 Build and manage S/390 microprocessor IOCDs and IPL attributes**.
9. A list of processors that are accessible from the system that HCD is running on are displayed. Refer to Figure 6-29. Type s next to the processor in our example.

```

S/390 Microprocessor Cluster List                               Row 1 of 6
Command ==> _____ Scroll ==> CSR

Select one or more CPCs, then press Enter.

-----CPC----- IODF
/ SNA Address      Type  Model Processor ID
# IBM390PS.H40
# IBM390PS.TC8M
_ USIBMSC.SCZP101  2094 S18   SCZP101
_ USIBMSC.SCZP201  2097 E26   SCZP201
s USIBMSC.SCZP301  2817 M32   SCZP301
_ USIBMSC.SCZP901  2084 C24   SCZP901

```

Figure 6-29 HCD Writing an IOCDs from an IODF

The next panel shows a summary of all IOCDSs for this processor. Here we can see the current IOCDS that is in use by this processor indicated by the YES-POR under the Write Protect column. Refer to Figure 6-30.

10.Type u next to IOCDS A3 to update this IOCDS with our new IODF information.

```

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

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

/ IOCDS      Name      Type      Status      -----Token Match----- Write
_ A0.SCZP301  IODF77    LPAR     Alternate   No          No          No
_ A1.SCZP301  IODF00    LPAR     Alternate   No          No          No
_ A2.SCZP301  IODF05    LPAR     POR         Yes         No          Yes-POR
u A3.SCZP301  IODF76    LPAR     Alternate   No          No          No
  
```

Figure 6-30 HCD IOCDS List

11.Press enter to proceed to the next panel. Refer to Figure 6-31.

12.Type IODF06 in the Title1 field. Leave the other options as defaults.

```

----- Build IOCDSs -----
Row 1 of 1
Command ===> _____ Scroll ===> CSR

Specify or revise the following values.

IODF name . . . . . : 'SYS6.IODF06'

Title1 . IODF06_____
Title2 : SYS6.IODF06 - 2011-11-20 11:41

IOCDS      Switch IOCDS      Write IOCDS in
A3.SCZP301 No          preparation of upgrade
No
  
```

Figure 6-31 HCD Build IOCDS

13.Press enter to continue to the final panel. See Figure 6-32.

```

----- Job Statement Information -----

Specify or revise the job statement information.

Job statement information
//IOCDSA3 JOB (ACCOUNT),'NAME',REGION=128M
//*
//*
//*
/*
  
```

Figure 6-32 HCD IOCDS Job Statement panel

14. Ensure the JCL statements meet your sites requirements and press enter to submit the Write IOCDS job.
15. After the job is completed, verify the system output and look for the following message in Example 6-5.

*Example 6-5 HCD IOCDS written successfully*

---

```
CBDA674I  IOCP successfully completed for A3.SCZP301.
```

---

16. Return to HCD Option 2.11, and navigate back to the IOCDS list for this processor, as shown in Figure 6-33.

Observe that IOCDS A3 now has a name of IODF06 and that there is now a Token Match of Yes under the IOCDS/Proc column.

|  |               | 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.SCZP301                                       | IODF77        | LPAR        | Alternate        | No        | No                          | No        |
| _ A1.SCZP301                                       | IODF00        | LPAR        | Alternate        | No        | No                          | No        |
| _ A2.SCZP301                                       | IODF05        | LPAR        | POR              | Yes       | No                          | Yes-POR   |
| _ <b>A3.SCZP301</b>                                | <b>IODF06</b> | <b>LPAR</b> | <b>Alternate</b> | <b>No</b> | <b>Yes</b>                  | <b>No</b> |

Figure 6-33 HCD IOCDS List after IODF update

## 6.5.2 Activating the production IODF

There are two ways to activate an IODF. One way is to use the HCD panels in Option 2 and another way is using the ACTIVATE system command. We demonstrate using the ACTIVATE system command:

1. Access a console or SDSF under TSO for this system where commands are authorized to be entered.
2. Enter **D IOS,CONFIG** to see what IODF the hardware and software are using for this system. Refer to Figure 6-34:
  - Active IODF data set: shows which IODF the Software (or LPAR) is currently using
  - Source: Shows which IODF the hardware (or HSA) for processor SCZP301 is currently using.

```
-D IOS,CONFIG
IOS506I 12.11.10 I/O CONFIG DATA 396
ACTIVE IODF DATA SET = SYS6.IODF05
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-16 16:39:06 SYS6   IODF05
ACTIVE CSS: 1 SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE
```

Figure 6-34 MVS IODF system display

3. Enter **ACTIVATE IODF=06,TEST** to test the new IODF, and verify that the changes we made to the IODF are reflected accurately in the test display. Refer to Figure 6-35.

This command tests changes being made to the Hardware or HSA.

```
-ACTIVATE IODF=06,TEST  
IOS500I ACTIVATE RESULTS 363  
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE  
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:  
      1.411C-1.411F,1.513C-1.513F,1.511C-1.511F,1.413C-1.413F  
      COMPID=SC1XL  
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR  
      SCZP301:  
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917  
      COMPID=SC1XL  
NOTE = A888,F  
      COMPID=SC1XL  
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR  
      SCZP301: 1.411C,1.513C,1.511C,1.413C  
      COMPID=SC1XL  
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR  
      SCZP301: 1.9000,1.9100  
      COMPID=SC1XL  
NOTE = A879,FOLLOWING CHANNEL PATHS ARE TO BE DELETED FROM PROCESSOR  
      SCZP301: 1.22,1.26,2.22,2.26,3.22,3.26,0.22,0.26  
      COMPID=SC1XL  
NOTE = A881,FOLLOWING CHANNEL PATHS ARE TO BE ADDED TO PROCESSOR  
      SCZP301: 0.22,0.26,1.22,1.26,2.22,2.26,3.22,3.26  
      COMPID=SC1XL  
NOTE = 0112,REQUEST CONTAINS DELETE(S), SPECIFY FORCE ON H/W ACTIVATE  
COMPID=SC1C3
```

Figure 6-35 MVS Test ACTIVATE display

4. Enter **ACTIVATE,IODF=06,SOFT,TEST** to test changes being made to the software or operating system (TEST2094) for this system. Refer to Figure 6-36.

```
-ACTIVATE IODF=06,SOFT,TEST  
IOS500I ACTIVATE RESULTS 431  
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE  
NOTE = 0100,SOFTWARE-ONLY CHANGE  
      COMPID=SC1C3  
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS WILL BE OUT OF SYNC  
      COMPID=SC1C3
```

Figure 6-36 MVS Test ACTIVATE SOFT command

The responses to the displays looks positive as per the message **TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE**.

We are now ready to activate the new IODF to the hardware and software for this system.

Because our system (SC30) is part of a sysplex (COMMPLEX), we must first activate the IODF on those systems. The systems defined in this sysplex are SC30, SC31, SC32, and SC33. Refer to Figure 6-37 on page 184.

```

-D XCF
IXC334I 10.18.08 DISPLAY XCF 305
  SYSPLEX COMPLEX:  SC30          SC31          SC32
                   SC33

```

Figure 6-37 MVS D XCF command

A SOFT ACTIVATE was performed already on systems SC32 and SC33.

- For system SC31, enter **ACTIVATE IODF=06,SOFT**.

```

-RO SC31,ACTIVATE IODF=06,SOFT
IOS1500I ACTIVATE RESULTS 887
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 0100,SOFTWARE-ONLY CHANGE
  COMPID=SC1C3
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS ARE NOW OUT OF SYNC
  COMPID=SC1C3

```

Figure 6-38 MVS ACTIVATE SOFT command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 6-38.

- Enter **D IOS,CONFIG** to verify that the activate for the LPAR SC31 completed.

```

-RO SC31,D IOS,CONFIG
IOS506I 12.20.50 I/O CONFIG DATA 893
ACTIVE IODF DATA SET = SYS6.IODF06
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-16 16:39:06 SYS6  IODF05
ACTIVE CSS: 1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 6-39 MVS D IOS.CONFIG command

Observe that the software is on **SYS6.IODF06** and the hardware is on **SYS6.IODF05**. Refer to Figure 6-39.

- For system SC30 enter **ACTIVATE IODF=06**.



```

-RO SC30,ACTIVATE IODF=06,FORCE
IOS1500I ACTIVATE RESULTS 511
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:
      1.411C-1.411F,1.513C-1.513F,1.511C-1.511F,1.413C-1.413F
      COMPID=SC1XL
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301:
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917
      COMPID=SC1XL
NOTE = A888,F
      COMPID=SC1XL
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 1.411C,1.513C,1.511C,1.413C
      COMPID=SC1XL
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301: 1.9000,1.9100
      COMPID=SC1XL
NOTE = A879,FOLLOWING CHANNEL PATHS ARE TO BE DELETED FROM PROCESSOR
      SCZP301: 1.22,1.26,2.22,2.26,3.22,3.26,0.22,0.26
      COMPID=SC1XL
NOTE = A881,FOLLOWING CHANNEL PATHS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 0.22,0.26,1.22,1.26,2.22,2.26,3.22,3.26
      COMPID=SC1XL

```

Figure 6-40 MVS ACTIVATE IODF command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 6-40.

**Note:** FORCE was specified in this ACTIVATE command as the CHPIDs in this IODF were previously defined to this processor (SCZP301), but we needed to add Switch information to them. This required the Activate process to take the Channels offline then online again.

8. Enter D IOS,CONFIG to verify that the activate for the LPAR SC30 completed.

```

-RO SC30,D IOS,CONFIG
IOS506I 12.27.04 I/O CONFIG DATA 526
ACTIVE IODF DATA SET = SYS6.IODF06
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-20 11:41:25 SYS6      IODF06
ACTIVE CSS: 1 SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 6-41 MVS D IOS,CONFIG command

Observe that the software is on SYS6.IODF06 and the hardware is now on SYS6.IODF06. Refer to Figure 6-41.

The IOCDS must be switched to IOCDS A3 so that the Reset Profile for this processor SCZP301 points to the correct IODF information should there be a scheduled or unscheduled POR of this processor.

9. Enter **ACTIVATE ACTIOCDS=A3** to switch the IOCDS.

```

-RO SC30,ACTIVATE ACTIOCDs=A3
IOS500I ACTIVATE RESULTS 538
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 010C,THE ACTIVE IOCDs HAS BEEN SWITCHED TO A3
COMPID=SC1C3

```

Figure 6-42 MVS ACTIVATE ACTIOCDs command

Observe the message THE ACTIVE IOCDs HAS NOW BEEN SWITCHED TO A3. Refer to Figure 6-42.

10. Return to HCD, and select:

- SYS6.IODF06
- Option 2.11
- Our test processor SCZP301

IOCDs A3 is now selected and a 'Yes' response is displayed for all columns. Refer to Figure 6-43.

```

IOCDs List Row 1 of 4 More: >
Command ==> _____ Scroll ==> CSR

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

-----Token Match----- Write
/ IOCDs Name Type Status IOCDs/HSA IOCDs/Proc. Protect
_ A0.SCZP301 IODF77 LPAR Alternate No No No
_ A1.SCZP301 IODF00 LPAR Alternate No No No
_ A2.SCZP301 IODF05 LPAR Alternate No No No
_ A3.SCZP301 IODF06 LPAR POR Yes Yes Yes-POR

```

Figure 6-43 HCD IOCDs List after new IODF activate

11. Press F20 to display information in the IOCDs List panel. See Figure 6-44.

```

IOCDs List Row 1 of 4 More: <
Command ==> _____ Scroll ==> CSR

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

--Last Update--
/ IOCDs Date Time IOCDs Configuration Token Information
_ A0.SCZP301 2011-11-07 13:46 SCZP301 15:05:04 11-11-03 SYS6 IODF77
_ A1.SCZP301 2011-11-16 17:53 SCZP301 09:25:30 11-11-10 SYS6 IODF00
_ A2.SCZP301 2011-11-16 18:12 SCZP301 16:39:06 11-11-16 SYS6 IODF05
_ A3.SCZP301 2011-11-20 12:09 SCZP301 11:41:25 11-11-20 SYS6 IODF06

```

Figure 6-44 HCD IOCDs List after pressing F20

Now that the hardware component (HSA) of the IODF is synchronized, we must make sure that the software component is also synchronized.

The volume where the IODF was saved before (IODFPK) also contains a dataset called SYS#.IPLPARM. This dataset contains LOADxx members that are used at IPL startup to

determine certain system parameters. It also tells the system being IPLed which IODF to use.

When an LPAR is IPLed, a Load (IPL) address and a Load parameter (loadparm) are specified.

Figure 6-45 shows what the Load Profile looks like for this system. In Figure 6-45:

- The LOAD address points to where z/OS and programs reside
- The LOAD parameter C73001M1:
  - **C730** equals where the IODF resides and also the SYS#.IPLPARM dataset
  - **01** equals the suffix value of the LOADxx member in SYS#.IPLPARM
  - **M** equals the NIP suppression indicator
  - **1** equals alternate nucleus member

|   |  |  |
|---|--|--|
| Profile name                              | TEST9302   |  |
| Description                               | Z1DRB1 GA  |  |
| Load type                                 | <input type="radio"/> Normal <input checked="" type="radio"/> Clear <input type="radio"/> SCSI <input type="radio"/> SCSI dump |  |
| Load address                              | 9302   | <input type="checkbox"/> Use dynamically changed address   |
| Load parameter                            | C73001M1   | <input type="checkbox"/> Use dynamically changed parameter |
| Time-out value                            | 300  | 60 to 600 seconds  |
| <input type="checkbox"/> Store status     |  |  |
| Worldwide port name                       | 0  |  |
| Logical unit number                       | 0  |  |
| Boot program selector                     | 0  |  |
| Boot record logical block address         | 0  |  |
| Operating system specific load parameters |  |  |

Figure 6-45 HMC Load Profile

12. We must logon to the system where the IODF was written, in our example SC30 and access SYS#.IPLPARM.

13. Enter the command **D IPLINFO** to obtain which LOADxx member the system used during IPL. Refer to Figure 6-46.

```

-D IPLINFO
IEE254I 14.34.15 IPLINFO DISPLAY 506
SYSTEM IPLLED AT 22.03.35 ON 11/02/2011
RELEASE z/OS 01.13.00 LICENSE = z/OS
USED LOAD01 IN SYS0.IPLPARM ON OC730
ARCHLVL = 2 MTLSHARE = N
IEASYM LIST = XX
IEASYS LIST = (00) (0P)
IODF DEVICE: ORIGINAL(OC730) CURRENT(OC730)
IPL DEVICE: ORIGINAL(O9302) CURRENT(O9302) VOLUME(Z1DRB1)

```

Figure 6-46 MVS D IPLINFO command

Here is a sample of the LOAD01 member from SYS0.IPLPARM showing:

- HWNAME = SCZP301
- LPARNAME = A11
- IODF = 06

See Example 6-6.

*Example 6-6 HCD - LOAD01 member in SYS0.IPLPARM*

---

```
*-----DEFINITION FOR SC30-----*
HNAME    SCZP301
LNAME    A11
*IEASYM   99
SYSPLEX  COMPLEX Y
IODF     06 SYS6      TEST2094 01 Y
SYSCAT   COMCAT123CMCAT.COMCAT
PARMLIB  SYS1.PARMLIB
PARMLIB  CPAC.ZOSR1D.PARMLIB
PARMLIB  SYS1.IBM.PARMLIB
*-----DEFINITION FOR SC31-----*
HNAME    SCZP301
LNAME    A13
*IEASYM   99
SYSPLEX  COMPLEX Y
IODF     06 SYS6      TEST2094 01 Y
SYSCAT   COMCAT123CMCAT.COMCAT
PARMLIB  SYS1.PARMLIB
PARMLIB  CPAC.ZOSR1D.PARMLIB
PARMLIB  SYS1.IBM.PARMLIB
```

---

This ensures that the operating system component of the IODF (software) when IPLed, matches the hardware component of the IODF (hardware) that is loaded into HSA.

## 6.6 Configuring the IBM Storage System DS8000

The host adapter ports in the DS8000 storage control unit must be configured for FICON (FC) protocol. Follow the procedure described in Appendix D, “Configuring the DS8000 for FICON” on page 415, to configure the host adapter ports.

For planning and implementation information, refer to *IBM System Storage DS8000: Architecture and Implementation*, SG24-8886.

### 6.6.1 Configuring the FICON Director

The FICON Director must be configured to match the designated configuration (shown in Figure 6-1 on page 156). For the desired switched topology, this includes:

- ▶ Specifying the Domain ID in the Director
- ▶ Enabling the ports
- ▶ Checking the port settings

Other settings in the FICON Director are optional and must be specified after reading the considerations explained in Chapter 4, “Planning the FICON environment” on page 77.

Make sure that the selected ports (LX or SX) in the Director match the FICON node (LX or SX) that will be attached.

All the required steps to configure the FICON Director are described in detail in Chapter 8, “Configuring FICON Directors” on page 245. Go to 8.1.1 “Configuration flowchart” on page 246 and follow the procedures described there.

Return to this section after the FICON Director is configured and ready for use.

## 6.6.2 Connecting the fiber optic cables

After all definition and configuration tasks are completed, the fiber optic cables must be plugged into the FICON channels, CU host adapters, and the FICON Director.

All of the ports in our configuration are long wavelength (LX) laser types. For LX type lasers, a 9µm single mode (SM) fiber optic cable is recommended to achieve maximum performance and distance. All fiber optic cables must have an LC duplex connector at both ends.

Plug the fiber optic cable connectors to the designated ports, shown in Figure 6-1 on page 156, to the FICON channel, the CU, and the FICON Director, as listed here:

- ▶ PCHID 39C to FICON Director port 00
- ▶ PCHID 3E4 to FICON Director port 19
- ▶ CU host adapter port 0033 to FICON Director port 34
- ▶ CU host adapter port 0233 to FICON Director port 62

Make sure the fiber optic cable ends are cleaned before plugging them into the transceiver. Ensure that all fiber optic cables are labeled and documented for later use (for example, when recabling or for problem determination).

For further information and considerations regarding fiber optic cabling and documentation, refer to 4.11 “Physical connectivity” on page 113 and 4.2 “Documentation” on page 80.

## 6.6.3 Configuring the channel path and device online

After you successfully activate the new I/O definitions, you can configure the channel path online and try to communicate with the devices.

Even if all of the definitions in the IOCDs are correct, the channel status might be offline.

To configure the channel path and device online:

1. Enter **D M=CHP(22)** on the system console to display the current status of CHPID 22. If the CHPID is offline, enter **CF CHP(22),ON** to configure the channel online. The channel status should change to ONLINE, as shown in Example 6-7.

*Example 6-7 MVS - CHPID 22 online state*

---

```
- CF CHP(22),ONLINE
IEE502I CHP(22),ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

---

To achieve the best performance on the FICON channel, make sure that zHPF is enabled. Refer to 4.9.6 “High Performance FICON” on page 107 for considerations regarding how to exploit zHPF on System z servers.

2. Enter **D IOS,ZHPF** at the z/OS console to display zHPF settings for the z/OS image. If zHPF is disabled, enter **SETIOS ZHPF,YES**. This enables zHPF, as shown in Example 6-8.

*Example 6-8 MVS - Displaying zHPF*

---

```
-D IOS,ZHPF
IOS630I 13.27.23 ZHPF FACILITY 021
HIGH PERFORMANCE FICON FACILITY IS ENABLED
```

---

Using the **SETIOS ZHPF=YES** command enables zHPF temporarily. However, after the next system IPL, the zHPF facility is reset to the default (disabled). To permanently enable zHPF for z/OS, add or change the ZHPF parameter in the SYS1.PARMLIB member IECIOSxx to ZHPF=YES. Refer to Example 6-9.

*Example 6-9 zHPF settings in SYS1.PARMLIB(IECIOS00)*

---

```

EDIT          SYS1.PARMLIB(IECIOS00) - 01.16          Columns 00001 00072
Command ==>                                         Scroll ==> CSR
***** ***** Top of Data *****
000001 HYPERPAV=YES
000002 ZHPF=YES
***** ***** Bottom of Data *****

```

---

- Now you can query the status and functional details of the channel by entering **D M=CHP(22)** at the operator console. The command output is shown in Example 6-10. It provides information about the channel and the attached devices.

*Example 6-10 MVS - D M=CHP(22)*

---

```

-D M=CHP(22)
IEE174I 18.24.03 DISPLAY M 163
CHPID 22: TYPE=1B, DESC=FICON SWITCHED, ONLINE
DEVICE STATUS FOR CHANNEL PATH 22
      0 1 2 3 4 5 6 7 8 9 A B C D E F
0413 . . . . . . . . . . . . + + + +
0900 + + + + + + + + + + + + + + + +
0901 + + + + + + + + + + + + + + + +
0902 + + + + + + + + + + + + + + + +
....
0910 + + + + + + + + + + + + + + + +
0911 + + + + + + + + + + + + + + + +
0912 + + + + + + + + + + + + + + + +
...
0917 + + * * * * * * * * * * * * * * +
1900 UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL UL
1901 UL UL UL UL UL UL HA HA HA HA HA HA HA HA HA HA
1902 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
...
1911 UL UL UL UL UL UL HA HA HA HA HA HA HA HA HA HA
1912 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1913 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1914 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1915 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1916 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1917 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
SWITCH DEVICE NUMBER = NONE
ATTACHED ND = SLKWRM.DCX.BRD.CA.1AFX0642C00Z
PHYSICAL CHANNEL ID = 039C
FACILITIES SUPPORTED = ZHPF

```

---

Example 6-10 on page 190 shows that CHPID 22 is online and operating in a switched configuration. Information is also displayed about the attached devices and the facilities (for example: ZHPF) supported by the channel.

4. To verify that communication to the attached devices works properly, enter **D M=DEV(xxxx)**, where xxxx is any device number. For example, to check the status of device 9000, which is a storage device, enter **D M=DEV(9000)** on a z/OS console, as shown in Example 6-11.

*Example 6-11 MVS - D M=DEV(9000) with device offline*

---

```

-D M=DEV(9000)
IEE174I 17.39.47 DISPLAY M 940
DEVICE 9000 STATUS=OFFLINE
CHP              40  41  42  43  20  24  22  26
ENTRY LINK ADDRESS 22  22  1E  1E  ..  ..  6500 6519
DEST LINK ADDRESS  1B  1B  10  10  0D  0D  6534 6562
PATH ONLINE       Y  Y  Y  Y  N  N  Y  Y
CHP PHYSICALLY ONLINE Y  Y  Y  Y  N  N  Y  Y
PATH OPERATIONAL  Y  Y  Y  Y  N  N  Y  Y
MANAGED          N  N  N  N  N  N  N  N
CU NUMBER        9000 9000 9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND        = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED    = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED   = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW

```

---

**Note:** Observe that the ZHPF function is not displayed for this device. We found in our tests that after the device was brought online, zHPF showed as enabled. Refer to Example 6-12.

*Example 6-12 MVS - D M=DEV(9000) with device online*

---

```

-D M=DEV(9000)
IEE174I 17.39.47 DISPLAY M 940
DEVICE 9000 STATUS=ONLINE
CHP              40  41  42  43  20  24  22  26
ENTRY LINK ADDRESS 22  22  1E  1E  ..  ..  6500 6519
DEST LINK ADDRESS  1B  1B  10  10  0D  0D  6534 6562
PATH ONLINE       Y  Y  Y  Y  N  N  Y  Y
CHP PHYSICALLY ONLINE Y  Y  Y  Y  N  N  Y  Y
PATH OPERATIONAL  Y  Y  Y  Y  N  N  Y  Y
MANAGED          N  N  N  N  N  N  N  N
CU NUMBER        9000 9000 9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND        = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED    = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED   = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW, ZHPF

```

---

The response to **D M=DEV(9000)** displays all available paths to the devices and their status. Information about the device (for example, device type), the control unit number, and the functions supported by the device (for example, MIDAW and ZHPF) is shown.

## 6.6.4 Verifying the installation

After all of the configuration tasks in the z196 server and the DS8000 storage are complete, you can verify that the current configuration matches the desired configuration shown in Figure 6-1 on page 156.

Perform this verification to ensure that the fiber optic cables are correctly plugged to the designated ports in the System z196 server and the DS8000 and to check the status of channels and devices:

1. Check the status of PCHID 39C in the z196 server. From any Hardware Management Console (HMC) where the desired CPC (server) is defined (for example, SCZP301), start a Single Object Operation (SOO) to that CPC's Support Element.

Display the list of channels installed on the CPC. Search the list for PCHID (39C) and double-click the PCHID icon to display the PCHID details panel shown in Figure 6-47.

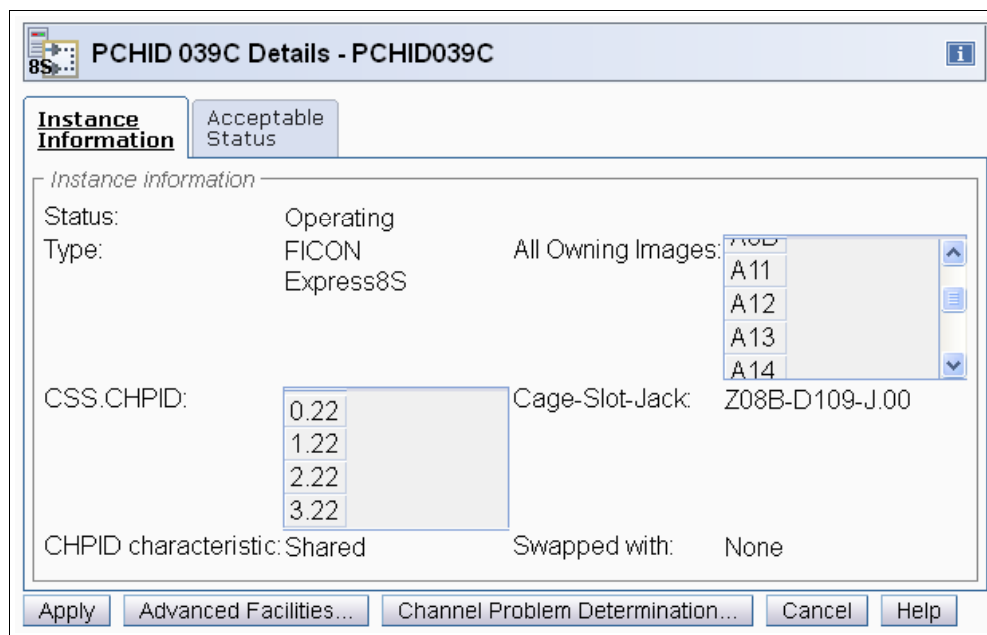


Figure 6-47 HMC PCHID details for 39C

Essential information for PCHID 39C is shown on the PCHID Details panel:

- The PCHID status is Operating
  - The PCHID type is FICON Express8S
  - CSS.CHPID 1.22 is assigned to PCHID 39C (plus others)
  - The owning images of PCHID 39C are A11 (plus others)
  - The CHPID assigned to PCHID 39C is shared
2. The CHPID Details panel provides information that is similar to the PCHID details panel. To display the CHPID details, select the channel list for an LPAR from the Support Element. Refer to Figure 6-48 on page 193.






Figure 6-48 HMC CHPID details for 1.22

Notice that the information for PCHID 39C (CHPID 1.22) provided on the details window matches the designated configuration. This proves that PCHID 39C has CHPID 1.22 assigned and that image A11 can access the channel. Repeat these checks on other channels that were recently configured.

**Important:** If any of the data displayed on the PCHID or CHPID detail panel does not match the desired configuration, you must correct the definitions in the IOCDs.

3. Check that the FICON channels are connected to the correct FICON switch and port. On either the PCHID details or CHPID details panel, click **Channel Problem Determination**. This will display the Channel Problem Determination panel, where you can select which information you want to display. See Figure 6-49.

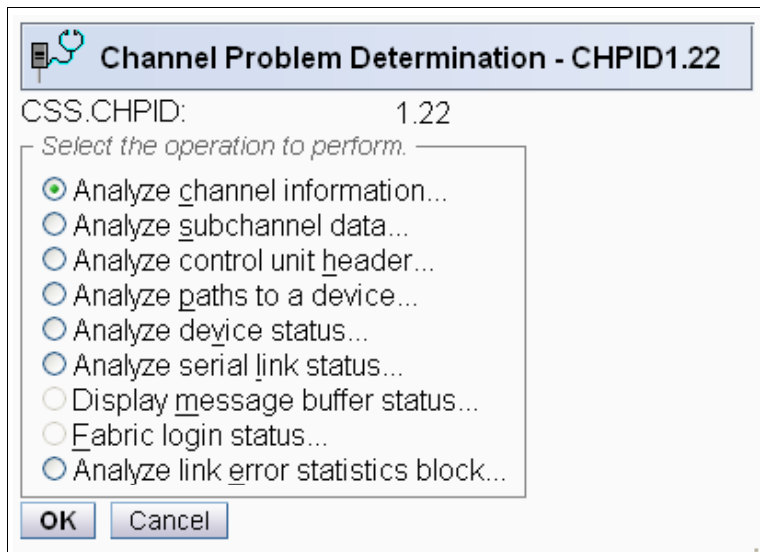


Figure 6-49 HMC Channel Problem Determination panel

4. Select **Analyze channel information**, and click **OK**. The Analyze Channel Information window is displayed, as shown in Figure 6-50, which provides information about the node attached to the FICON channel.

| Analyze Channel Information - CHPID1.22  |              |   |                           |
|--|--------------|---|---------------------------|
| Channel type:  | FICON Native | Hardware type:                            | E0                        |
| Partition ID:  | 11           | Hardware subtype:                         | 04                        |
| MIF image ID:  | 1            | 2 byte control unit link address defined: | Yes                       |
| Channel mode:  | Spanned      | Absolute address:                         | 000000009F4E9400          |
| CHPARM:  | 00           |   |                           |
| CSS.CHPID:   | 1.22         |   |                           |
| PCHID:   | 039C         |   |                           |
| Switch number:   | 65           |   |                           |
| Switch number valid:   | 1            | IFCC threshold:                           | 4                         |
| State:   | Online       | Channel link address:                     | 650000                    |
| Status:  | Operating    | Temp error threshold:                     | 4                         |
| Image chnl state:  | Online       | Suppress:                                 | 0000000000000000          |
| Image chnl status:   | Operating    | SAP Affinity:                             | 05                        |
| Error code:  | 00           | Card description:                         | FICON Express8<br>10km LX |
| Ber inbound:   | 0            | Connection rate:                          | 8Gbps                     |
| Ber outbound:  | 0            |   |                           |
| Node type:   | Self         | Node type:                                | Attached                  |
| Node status:   | Valid        | Node status:                              | Valid                     |
| Flag/parm:   | 10000722     | Flag/parm:                                | 00200A00                  |
| Type/model:  | 002817-M32   | Type/model:                               | SLKWRM-DCX                |
| MFG:   | IBM          | MFG:                                      | BRD                       |
| Plant:   | 02           | Plant:                                    | CA                        |
| Seq. number:   | 0000000B3BD5 | Seq. number:                              | 1AFX0642C00Z              |
| Tag:   | F022         | Tag:                                      | 6500                      |
| World wide node name: 5005076400CB3BD5   |              | World wide node name: 100000051E439201    |                           |
| World wide port name: 5005076028BF0B41   |              | World wide port name: 200000051E439201    |                           |
| <input type="button" value="OK"/> <input type="button" value="Error Details..."/> <input type="button" value="Refresh"/> |              |   |                           |

Figure 6-50 HMC Analyze Channel Information

Information about the nodes is displayed at the bottom part of the window. The lower left side displays information about the node in the z196 server. The lower-right side displays information about the attached node.

**Important:** Make sure that the Node status for both nodes is displayed as **Valid**. If any other status is shown, none of the displayed information is valid.

5. Check that the Type/model information and the serial number (Seq. number) is as expected.
6. Examine the Tag field for each node. The Tag provides information about the port number of the attached node. For the z196 server, the two right-most digits of the Tag value represent the CHPID number for the channel node (22). For the attached node, the four

digits represent the switch port number (6500). Be aware, however, that the tag value is provided by the attached device during link initialization and might have different meanings, depending on the vendor.

- The World Wide Node Name (WWNN) and the World Wide Port Name (WWPN) are also shown for each port, and can be used to prove that the channel is connected to the correct FICON adapter port if the WWNN or WWPN of the attached device is known.

If the node status is not Valid or the Tag value and WWPN value are not correct, check the fiber optic cable link between the z196 server and the FICON Director or DS8000 host adapter to ensure that it is plugged correctly.

In our scenario, we are now sure that PCHID 39C has connectivity to switch 65 and port 00, which matches our desired configuration. See Figure 6-1 on page 156.

If the displayed values are not as expected, the fiber optic cables might not be plugged correctly and must be checked.

After completing the preceding steps and proving that the physical path to the DS8000 storage controller and the logical definitions of the link are correct, check that the path to the control unit image is initialized correctly and properly defined.

- On the Channel Problem Determination panel, shown in Figure 6-49 on page 193, select **Analyze Serial Link Status**, and click **OK**. The Analyze Serial Link Status window is displayed, as shown in Figure 6-51.

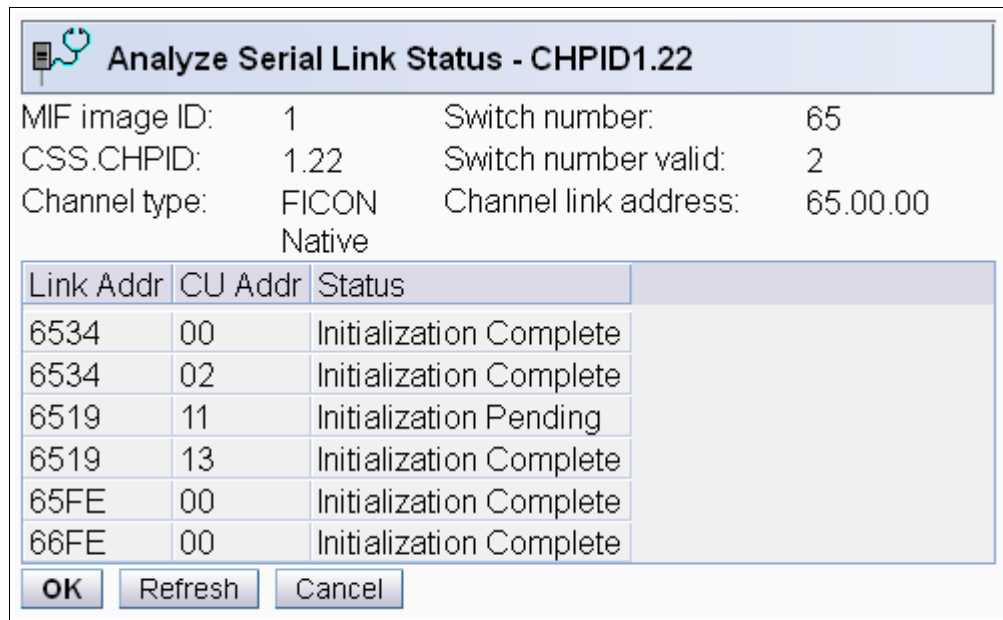


Figure 6-51 HMC Analyze Serial Link status

The Analyze Serial Link Status window provides status information about the link to the control unit images defined in the IOCDs. Scroll through the list of CU images and check that status for all CUs is Initialization Complete. In this case, Figure 6-51 shows a link status of Initialization Complete for all defined CU images on link address 6534.

Link address 6534 is the destination port where the DS8000 control unit port is physically attached. The link to CU address 00 and 02 is initialized. CU address 00 and 02 correspond to control unit number 9000 (CUADD=0) and 9100 (CUADD=2), as defined in the IODF.

Link address 6519 is the destination port in the FICON Director for the FCTC connection. CU address 13 represents the Partition Number for LPAR A13.

Link address 65FE is the CUP port in the FICON Director, which is used for communication and managing the FICON Director.

If the link status `Initialization Complete` is not shown, you must check that the ports in the FICON Director and CU are correctly configured and that the fiber optic cable link between the Director and the CU has the correct cable type and plugging.



## Configuring a cascaded topology

In this chapter, we describe all the tasks that are required to configure and define a FICON environment for a cascaded topology.

Establishing a cascaded topology involves:

- ▶ Describing our environment
- ▶ Using tasks and checklists:
  - Tasks
  - Verification checklist
- ▶ Defining the new configuration:
  - Defining the channel, CU, and storage devices
  - Defining our configuration example using HCD
  - Defining our configuration example using HCM
  - How the definition looks as IOCP statements
  - Defining a FICON Channel-to-Channel (FCTC) connection
- ▶ Building and activating the IODF:
  - Building the production IODF
  - Activating the production IODF
- ▶ Configuring the IBM Storage System DS8000:
  - Configuring the FICON Directors
  - Connecting the fiber optic cables
  - Configuring the channel path and device online
  - Verifying the installation

## 7.1 Establishing a cascaded topology

In this section, we describe the tasks required to configure a cascaded topology on a zEnterprise 196 server, including setting up of the FICON Directors and configuring the attached DS8000 storage subsystem. Our example is based on an infrastructure that is already built and installed with the following items physically installed:

- ▶ zEnterprise 196 server
- ▶ FICON Directors
- ▶ DS8000 storage subsystem
- ▶ Fiber cabling infrastructure

## 7.2 Description of our environment

Based on the considerations explained in Chapter 4, “Planning the FICON environment” on page 77, a cascaded FICON configuration is built, consisting of an IBM zEnterprise z196 server, an IBM SAN768B-2 and SAN384B-2 FICON Director, and an IBM DS8000 storage subsystem.

Longwave laser optic and SM fiber optic cables are used to achieve maximum distance and performance.

In our example case, the entire infrastructure is already installed in our environment. Our configuration is shown in Figure 7-1.

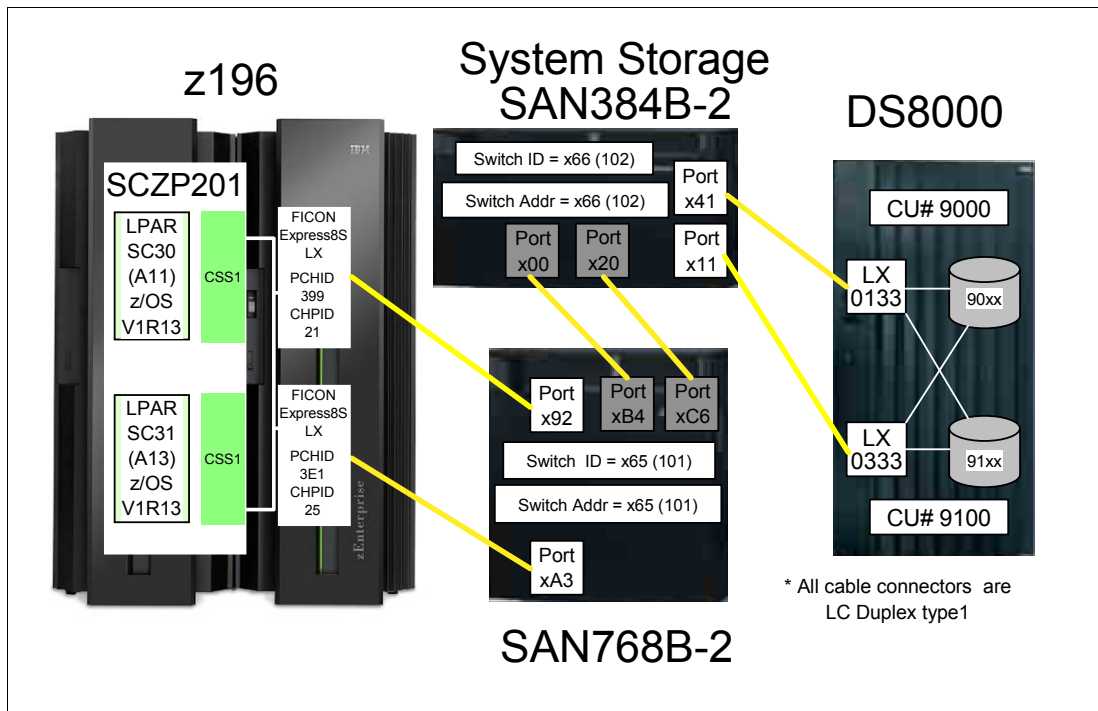


Figure 7-1 FICON Cascaded configuration

The z196 server (SCZP301) has two LPARs (A11 and A13) defined and activated. The system names running in this partition are SC30 and SC31, respectively. The operating system that is running in both partitions is z/OS V1R13. As a standard feature, the z196 server supports zHPF protocol, which is used for data transfer to the DS8000 storage device.

The two FICON channels are defined as spanned across CSS0, CSS1, CSS2, and CSS3 (only CSS1 is shown in our example). The channels are FICON Express8S LX features (FC 0409), using PCHID 399 (CHPID 21) and PCHID 3E1 (CHPID 25).

The SAN768B-2 FICON Director has connectivity to the z196 server on port x92 and port xA3. All the ports in the FICON Director are longwave (LX) ports. The SwitchID (Switch #) and the Switch Address (Switch @) are both set to (hex) 65. These are specified in the IODF. The Switch Address is the Domain ID specified in the FICON Director and is defined as decimal, for example:

SwitchID / Switch Address of 65 (IODF) would equal a Domain ID of 101 (switch)

The SAN384B-2 FICON Director has connectivity to storage devices on port x41 and x11. All of the ports in the FICON Director are longwave (LX) ports. The Switch ID (Switch #) and the Switch Address (Switch @) are both set to (hex) 66. The Switch ID is specified in the IODF. The Switch Address is the Domain ID specified in the FICON Director.

The FICON Directors are connected by two Inter-Switch Link (ISL), which consist of two fiber optic cables connected to ports xB4 and xC6 in switch 65 and ports x00 and x20 in switch 66. If the FICON Directors are connected, they build a High Integrity Fabric, which is required to support cascaded FICON Directors. The ISLs are transparent in the path from the System z server to the CU and do not require any definitions in the IODF.

An IBM System Storage DS8000 is connected to the SAN384B-2 FICON Director on port 41 and the 11. The two host adapters are installed in the DS8000 supporting longwave (LX) laser. Port numbers 0133 and 0333 in the DS8000 are used to connect to the FICON Director. Two logical control units (9000 and 9100) are defined that have devices 9000-90FF or 9100-91FF assigned. The DS8000 has the zHPF feature enabled.

A maximum unrepeated distance of 10km (6.2 miles) is supported by the longwave laser (LX) feature at 16Gbps when using 9µm single mode (SM) fiber optic cables. The fiber optic cables have an LC Duplex connector at both ends to connect to the z196 FICON Express8S channels, to the ports in the FICON Directors, and to the DS8000 host adapters.

**Note:** The SAN384B-2 and SAN768B-2 FICON Directors support 25km ISL links with FC2881 at 8Gbps LX.

## 7.3 Tasks and checklist

In this section, we provide an overview of the required tasks and the checklist that we used to make sure that all hardware and software prerequisites were met before starting the configuration tasks.

### 7.3.1 Tasks

Figure 7-2 on page 200 shows the main steps that are required to define and activate a cascaded FICON Director configuration.

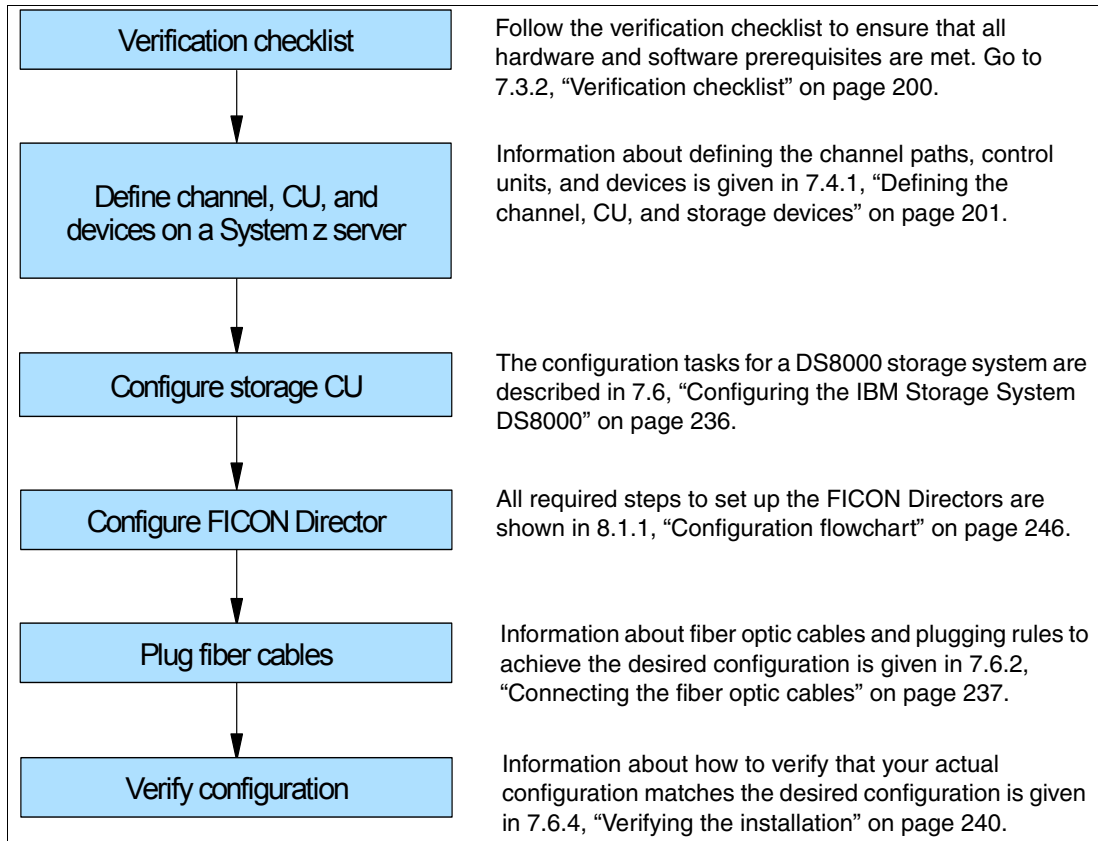


Figure 7-2 Steps for configuring and verifying a FICON cascaded configuration

### 7.3.2 Verification checklist

Before configuring the cascaded FICON topology shown in Figure 7-1 on page 198, the following list was checked. All steps in the checklist must be finished and corrected (if required) to ensure a smooth and successful configuration of the topology. Both hardware and software requirements must be checked.

- Check that FICON features are available on the System z server to establish the desired configuration. For details about each feature code, see "System z FICON feature support" on page 37.
  - FICON Express4
    - LX FC 3321, FC 3324
    - SX FC 3322
  - FICON Express8
    - LX FC 3325
    - SX FC 3326
  - FICON Express8S
    - LX FC 0409
    - SX FC 0410
- If using the FICON Express8S feature, check the System z operating system requirements.



- Check the 2817DEVICE PSP bucket for the latest information about FICON Express8S support on the operating system.
- Check that the number and types of FICON Director ports match the configuration requirements.
  - LX SFPs
  - SX SFPs
- Check DS8000 storage hardware requirements to support FICON longwave (LX) connectivity to the z10 server.
  - Check that FC 0709 and FC 7092 are installed to support zHPF.

**Note:** FC 0709 and FC 7092 must be ordered to obtain the license key to enable zHPF support in a DS8000 storage controller.

- Check that the DS8000 firmware is at level 6.1.16.0 or higher to support the zHPF feature.
- Check that the zHPF feature is enabled in the DS8000.
- Check that the correct fiber optic cables are available to connect the z196 server to the FICON Director, the two ISL links between the FICON Directors and the FICON Director to the DS8000 storage controller.
  - A 9µm single mode (SM) fiber optic cable is required to support longwave laser (LX) for a maximum distance of 10km (6.2 miles) at maximum speed.
  - A 50µm or 62.5µm multi mode (MM) fiber optic cable is required to support shortwave laser (SX). See “System z FICON feature support” on page 37 for the maximum supported distance, depending on cable type and speed.
  - An LC duplex connector is required at both ends of the fiber optic cables to connect to the z196 server FICON channel, the ports in the FICON Director, and the DS8000 host adapter ports.

**Note:** All fiber optic cables used in a link must be of the same type, for example, they must be either all single mode or all multi-mode fiber optic cables.

## 7.4 Defining the new configuration

In this section, we describe all of the tasks that are required to achieve the designated configuration shown in Figure 7-1 on page 198. We explain what must be done to get the FICON channels, FICON Director, storage control units, and the storage device online and operating.

For our scenario, we had an active partition in a zEnterprise 196 server running z/OS V1.13 with HCD. We used HCD to create, save, and activate our I/O definitions.

We also showed the definition process using HCM.

### 7.4.1 Defining the channel, CU, and storage devices

There are several tools available that you can use to define and configure a FICON environment on a System z server. The definition tools and an explanation of the statements and keywords used in the IOCP are described in Appendix C, “Configuration and definition

tools” on page 409. Only statements and keywords relevant to our desired configuration are explained. Refer to the *IOCP User's Guide*, SB10-7037, for details about all IOCP statements and keywords.

Steps using both HCD and HCM are used to illustrate the definition process for the CHPIDs, CUs, and devices for the test configuration we describe in this section. Refer to *HCD User's Guide*, SC33-7988, or the *HCM User's Guide*, SC33-7989, for more detailed information and procedural example.

After all definition tasks are completed, the definition data is stored in the I/O definition file (IODF) data set. Any time you make changes to the I/O configuration in the “Work” IODF, a new “Production” IODF must be built and saved to the IODF volume (software) along with updating and writing to a spare IOCDS (hardware). This is required to maintain synchronization of the IODF between the software and hardware definitions.

The newly created “Production” IODF is used to write IOCDS information to the Support Element of the System z server. The Support Element in the System z server stores the IOCDS on its built-in HDD. The IOCDS stored in the Support Element is loaded into the server's storage (HSA) during a Power-on Reset (POR) of the System z server.

Changes or additions of channel paths, CUs, or devices can be performed dynamically in most cases. Because no system outage is required for dynamic changes, this is the preferred method for changing I/O definitions.

If changes to the I/O definitions are performed dynamically, those changes are directly updated into the Hardware System Area (HSA) of the server. The HSA is a separate storage area in the System z196 server that holds all I/O specifications and provides status information for all defined I/Os. The operating system uses the HSA to manage the I/O subsystem. For more details about System z channel subsystem characteristics, refer to 4.8.1, “System z” on page 97.

## 7.4.2 Defining our configuration example using HCD

To use HCD to define the configuration example:

1. Select the current work IODF using HCD. We use SYS6.IODF07.WORK for our example.
2. From the HCD main menu, select **Option 1.3 Processor list**.
3. Type `s` next to the processor ID to define/verify the CHPID definitions (for example SCZP301).
4. Type `s` next to the CSS ID, which contains the LPARs (for example A11 and A13).
5. Type `c` next to our CHPIDs to Change Channel Path Definition (CHPID 21). Refer to Figure 7-3 on page 203.

```

----- Change Channel Path Definition -----

Specify or revise the following values.

Processor ID . . . . : SCZP301
Configuration mode . : LPAR
Channel Subsystem ID : 1

Channel path ID . . . . 21   +           PCHID . . . 399
Channel path type . . . FC   +
Operation mode . . . . . SPAN +
Managed . . . . . No (Yes or No)   I/O Cluster _____ +
Description . . . . . Express8s LX

Specify the following values only if connected to a switch:

Dynamic entry switch ID __ + (00 - FF)
Entry switch ID . . . . __ +
Entry port . . . . . __ +

```

Figure 7-3 HCD Change Channel Path Definition

In Figure 7-3:

- Processor ID = SCZP301
- CSS = 1
- CHPID = 21
- PCHID = 399

6. Press enter, and observe the Define Access List panel. Check that the required LPARs have access to CHPID 22. Refer to Figure 7-4.

```

----- Define Access List -----
Row 19 of 43
Command ==> _____ Scroll ==> CSR

Select one or more partitions for inclusion in the access list.

Channel subsystem ID : 1
Channel path ID . . . : 21   Channel path type . . : FC
Operation mode . . . : SPAN   Number of CHPIDs . . : 1

/ CSS ID Partition Name   Number Usage Description
/ 1   A11                1    OS   COMPLEX SC30
/ 1    A12                 2     OS    VMLINUX9
/ 1   A13                3    OS   COMPLEX SC31
/ 1    A14                 4     OS
/ 1    A15                 5     OS
/ 1    A16                 6     OS    COMPLEX SC32
/ 1    A17                 7     OS    VMLINUX4
/ 1    A18                 8     OS    COMPLEX SC33
/ 1    A19                 9     OS    ITS0SSI6
/ 2    A2A                 A     OS    ITS0SSI3

```

Figure 7-4 HCD Define Access List

7. Repeat these steps for CHPID 25.
8. Return to the HCD main menu, and select **Option 1.2 Switches**.
9. Enter Add to create a new switch definition.
10. Enter the highlighted values as shown in Figure 7-5:
  - Switch ID = 66
  - Switch type = 2032
  - Switch address = 66
  - Installed port range = 00 - FF
  - Switch CU number = 0066
  - Switch device number = 0066

```

----- Add Switch -----

Specify or revise the following values.

Switch ID . . . . . 66 (00-FF)
Switch type . . . . . 2032 _____ +
Serial number . . . . . _____
Description . . . . . SAN384B-2 _____
Switch address . . . . . 66 (00-FF) for a FICON switch

Specify the port range to be installed only if a larger range
than the minimum is desired.

Installed port range . . 00 - FF +

Specify either numbers of existing control unit and device, or
numbers for new control unit and device to be added.

Switch CU number(s) . . . 0066 _____ +
Switch device number(s) . 0066 _____

```

Figure 7-5 HCD Add Switch

11. Press enter to return to the Switch List.
12. Return to the HCD main menu, and select **Option 1.3 Processors**.
13. Navigate to CHPID 21 in the Channel Path List.
14. Enter the highlighted values as follows, and refer to Figure 7-6 on page 205:
  - Dyn SW = 65
  - Entry SW = 65
  - Entry Port = 92

```

Channel Path List      Row 23 of 123 More:  >
Command ==>> _____ Scroll ==>> CSR

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

Processor ID . . . . : SCZP301
Configuration mode . : LPAR
Channel Subsystem ID : 1

      PCHID          Dyn Entry +
/ CHPID AID/P Type+ Mode+ Sw+ Sw Port Con Mng Description
- 20    398    FC   SPAN  ___ ___ ___      No Express8s LX
- 21    399    FC   SPAN  65 65 92      No Express8s LX
- 22    39C    FC   SPAN  65 65 00      No Express8s LX
- 23    39D    FC   SPAN  ___ ___ ___      No Express8s LX
- 24    3E0    FC   SPAN  ___ ___ ___      No Express8s LX
- 25    3E1    FC   SPAN  ___ ___ ___      No Express8s LX
- 26    3E4    FC   SPAN  65 65 19      No Express8s LX
- 27    3E5    FC   SPAN  ___ ___ ___      No Express8s LX

```

Figure 7-6 HCD Define Switch Ports to Channel Path

15. Perform the same for our other test CHPID 25.
16. Return to the HCD main menu, and select **Option 1.4 Control units**.
17. Scroll down the list until CU 9000 is located.
18. Type c next to CU 9000 to Change Control Unit Definition. See Figure 7-7.
19. In our example, this Control Unit already contains existing definitions. We add connections to Switch #66, CHPIDs 21 and 25, to these definitions.
20. Enter the highlighted values as shown in Figure 7-7:
  - Connected to switches = 66 66
  - Ports = 41 11

```

----- Change Control Unit Definition -----

Specify or revise the following values.

Control unit number . . . . 9000 +
Control unit type . . . . . 2107      +
Serial number . . . . . 0L30002421 +
Description . . . . . YU CU 0

Connected to switches . . . 61 62 63 64 65 65 66 66 +
Ports . . . . . 1B 1B 10 10 34 62 41 11 +

Define more than eight ports . . 2 1. Yes
                                   2. No

```

Figure 7-7 HCD Change Control Unit Definition

21. Press enter to go to the Processor / CU panel.

22. Enter the highlighted values:

- Channel Path ID. Link Address = 21.6641 and 25.6611, as shown in Figure 7-8.

```

Select Processor / CU      Row 7 of 23 More:      >
Command ==> _____ Scroll ==> CSR

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

Control unit number . . : 9000      Control unit type . . . : 2107

-----Channel Path ID . Link Address + -----
/ Proc.CSSID 1----- 2----- 3----- 4----- 5----- 6----- 7----- 8-----
_ SCZP301.0 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____
_ SCZP301.1 40.1B  41.1B  42.10  43.10  22.6534 26.6562 21.6641 25.6611
_ SCZP301.2 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____
_ SCZP301.3 40.1B  41.1B  42.10  43.10  _____ _____ _____ _____
  
```

Figure 7-8 HCD Adding CHPIDs 22 and 26 to SCZP301.1, CU 9000

23. Press Enter to observe the Modify Device Parameters panel. See Figure 7-9.

```

Select Processor / CU
----- Modify Device Parameters -----
Row 1 of 3 More:      >
Command ==> _____ Scroll ==> CSR

Specify or revise any changes to the device parameters in the list below.
To view attached control units, scroll to the right.

Processor ID . . . . : SCZP301
Channel Subsystem ID : 1

-----Device-----      ---UA-----      Preferred Exposure
No., Range Type          SS+ Old New + Time-Out STADET CHPID + Device
9000,016 3390B          _  00 00  No   Yes   ___
9000,128 3390A          1  80 80  No   Yes   ___
9010,112 3390B          _  10 10  No   Yes   ___
***** Bottom of data *****
  
```

Figure 7-9 HCD Modify Device Parameters

This panel will display the Unit Addresses (UA) already defined for the device:

- Devices 9000-907F as 3390B (SubChannelSet or SS=0)
- Devices 9000-907F as 3390A (SubChannelSet or SS=1)

**Note:** With the addition of HyperPAV, we do not need to define 128 alias devices. Refer to 4.9.4, "Parallel Access Volume and HyperPAV" on page 105 to calculate your alias requirements.

24. Return to the HCD main menu, and select **Option 1.1 Operating system configurations.**

25. Type u next to the Operating system for our example (TEST2094).

26. Scroll down the list until I/O Device 9000 is located.
27. Type c next to 9000, and press enter until the Define Device to Operating System Configuration panel is displayed.
28. Type s next to TEST2094 to start the Device Parameters / Features panel. Refer to Figure 7-10.

```

----- Define Device to Operating System Configuration -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Select OSs to connect or disconnect devices, then press Enter.

Device number . : 9000          Number of devices : 16
Device type   . . : 3390B

/ Config. ID  Type   SS Description          Defined
_ ALLDEV     MVS    All devices             Yes
_ LABSERV1   MVS    Lab Services            Yes
_ L06RMVS1   MVS    Sysplex systems        Yes
_ MVS1       MVS    Production systems     Yes
_ OPENMVS1   MVS    OpenEdition MVS        Yes
s TEST2094   MVS    Sysplex systems        Yes

```

Figure 7-10 HCD Defining a Device to an Operating System Configuration

Here we can enable/disable the device-specific features by over typing the yes/no values. Refer to Figure 7-11.

```

----- Define Device Parameters / Features -----
                                                    Row 1 of 6
Command ==> _____ Scroll ==> CSR

Specify or revise the values below.

Configuration ID . : TEST2094    Sysplex systems
Device number   . . : 9000       Number of devices : 16
Device type     . . . : 3390B

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 *****

```

Figure 7-11 HCD Define Device Parameters / Features

29. Press Enter to go to the Assign/Unassign Device to Esoteric panel. Refer to Figure 7-12 on page 208.

```

----- Assign/Unassign Device to Esoteric -----
                                                    Row 1 of 9
Command ==> _____ Scroll ==> CSR

Specify Yes to assign or No to unassign.  To view devices already
assigned to esoteric, select and press Enter.

Configuration ID : TEST2094           Sysplex systems
Device number   . : 9000             Number of devices  : 16
Device type    . . : 3390B          Generic . . . . . : 3390

+-oteric Assigned Starting Number Number of Devices
- 01.ATL3      No      _____ _____
- 01.CART      No      _____ _____
- 01.NOVIO     Yes    _____ _____
- 01.SYSDA     Yes    _____ _____
- 01.SYSSQ     No      _____ _____
- 01.TAPE      No      _____ _____
- 01.TAPE3592 No      _____ _____
- 01.VIO       Yes    _____ _____

```

Figure 7-12 HCD Assign/Unassign Device to Esoteric

This panel allows us to assign devices to particular esoterics like SYSDA and VIO

30. Perform the same steps for the remainder of the devices for this control unit (3390B 9010-7F and 3390A 9000-7F), then return to the HCD main menu.

Steps 16 through 30 can now be performed for Control Unit 9100 and Devices 9100-917F

### 7.4.3 Defining our configuration example using HCM

To define the configuration example using HCM:

1. Select the current work IODF using HCM. We use SYS6.IODF07.WORK for our example. Refer to Figure 7-13 on page 209.






Figure 7-13 HCM File Open IODF06.hcm

2. Select **Locate** → **Processor**.
3. Double-click the required processor in the list (SCZP301). See Figure 7-14.




Figure 7-14 HCM Processor display

Our configuration already contains existing definitions, so we are adding to these.

4. Double-click the top pane of the processor display to edit the processor, or select **Edit** → **Processor**, and select the processor from the list. Refer to Figure 7-15 on page 210.

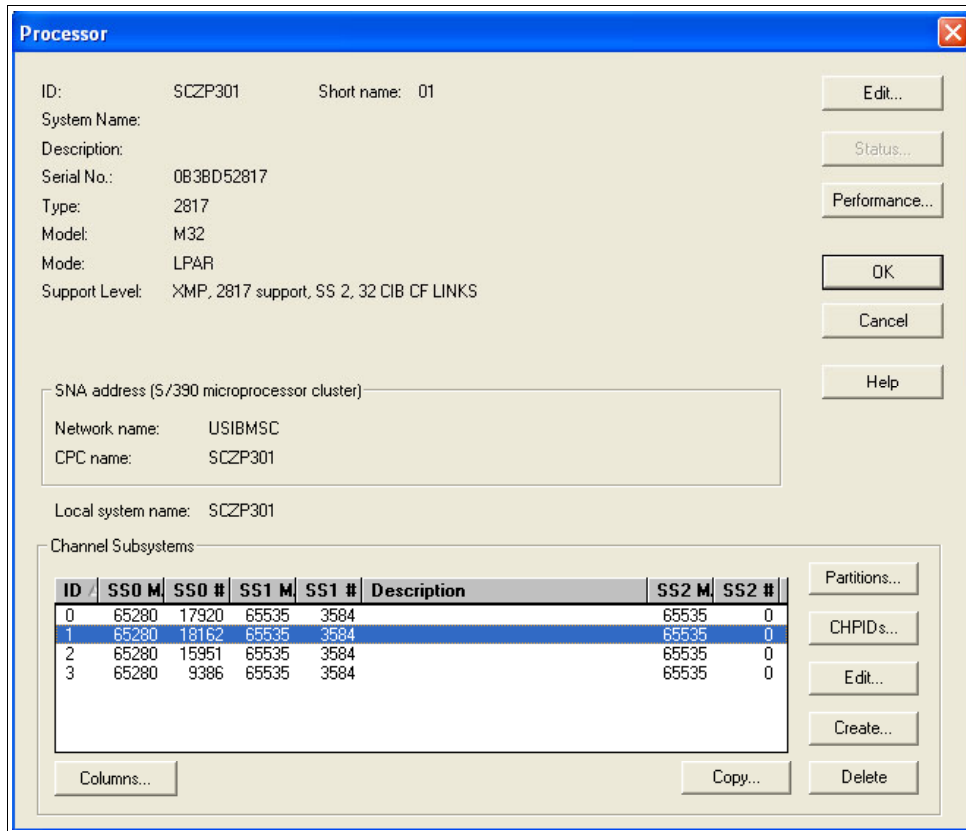


Figure 7-15 HCM Edit Processor

5. Highlight CSS ID 1, and click **CHPIDs**.
6. Scroll down the list until CHPIDs 21 and CHPIDs 25 are located. See Figure 7-16.

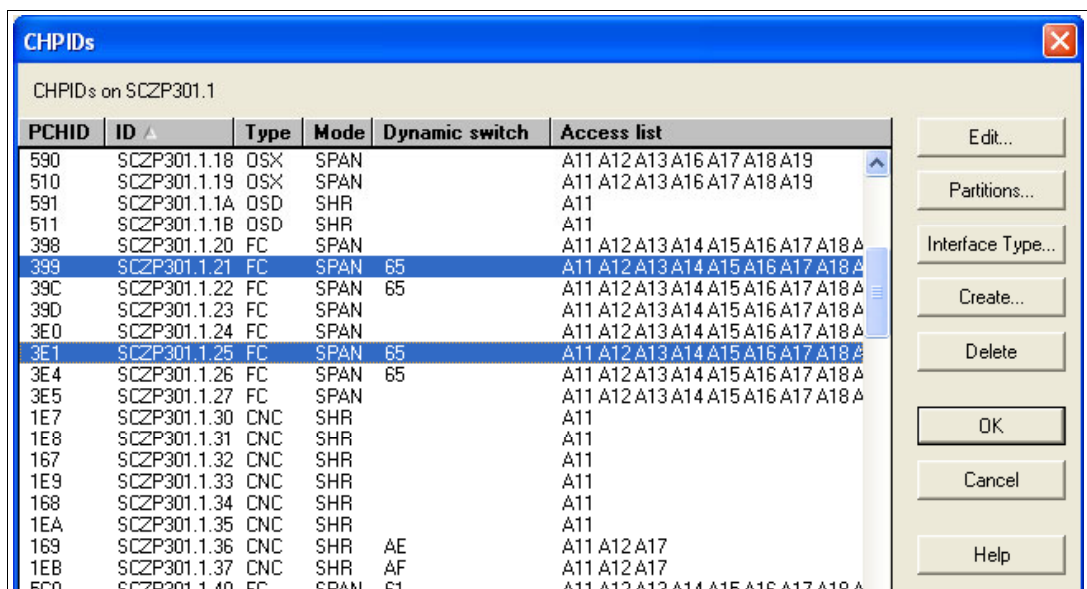


Figure 7-16 HCM Edit CHPIDs

7. Verify that they are defined correctly and have access to our test LPARs A11 and A13.

8. Select **Create** → **Switch**, and enter the following values, as shown in Figure 7-17:
- Switch ID = 66
  - Switch Address = 66
  - Type-Model = 2032
  - Installed Ports = 00 to FF
  - Control Unit Number = 0066
  - Device Number = 0066

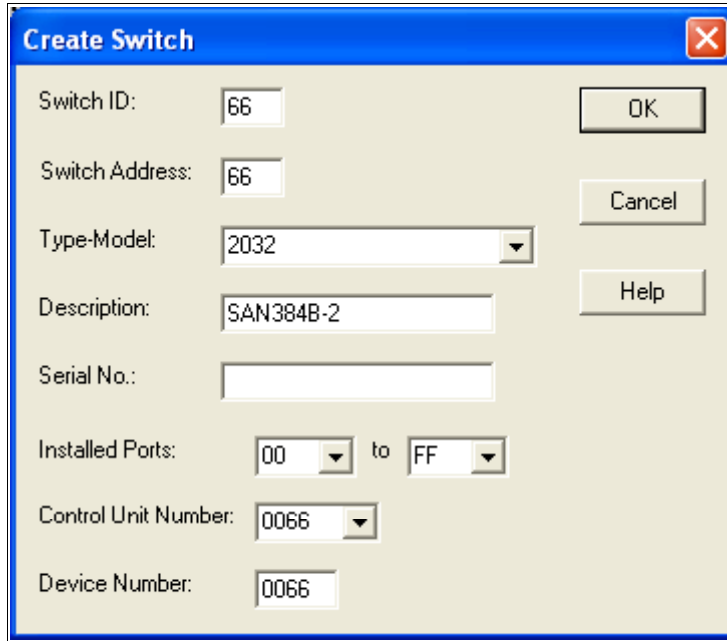


Figure 7-17 HCM Create Switch

9. Click **OK**.
10. Double-click the Switch #66 object to list all the defined ports.

**Note:** Because this Switch is connected to another Switch using Inter Switch Links (ISLs) and becomes cascaded, we must first connect the ISL ports between the Switches #65 and #66.

11. Highlight Port 00 on Switch #66, and then click **Connect**.
12. Navigate through the processors, CHPIDs, and Switches until you find Switch #65.
13. Expand the ports, highlight port B4, and then click **Connect** → **Physical Only**. Refer to Figure 7-18 on page 212.

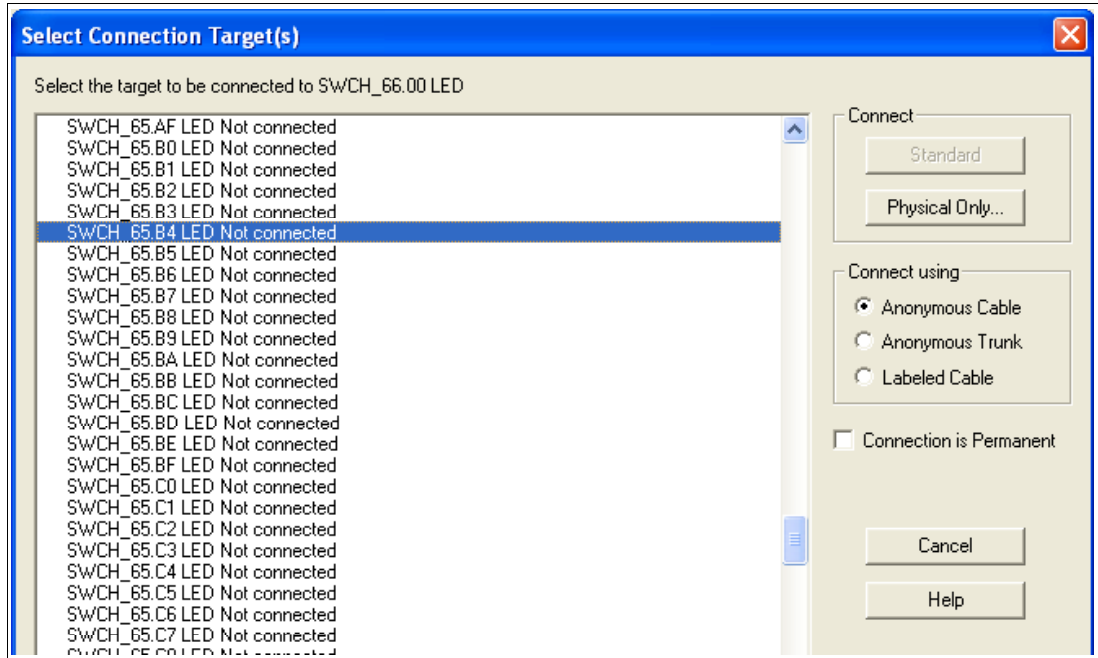


Figure 7-18 HCM Connect ISL ports

14. Perform the same action for Switch #66 port 20 and Switch #65 port C6. Refer to Figure 7-19.

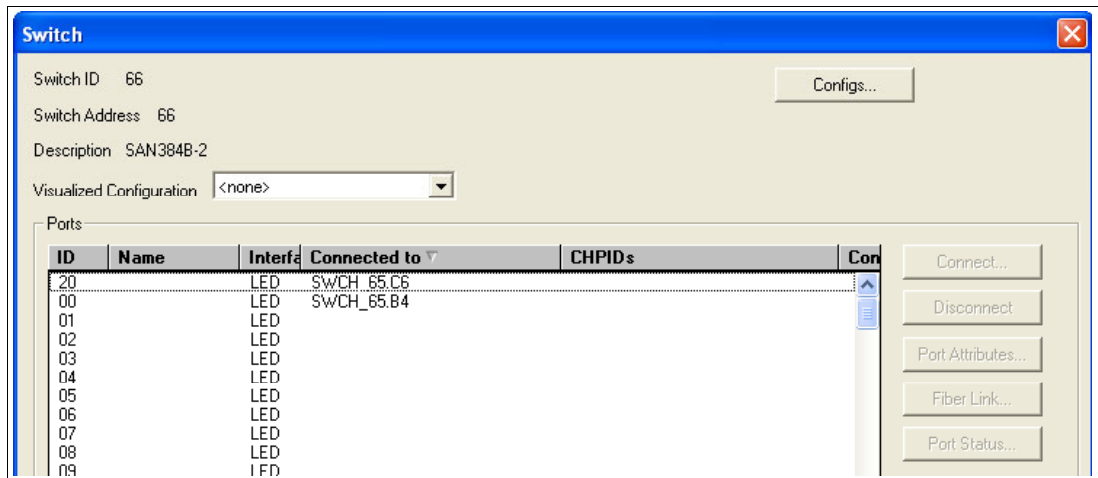


Figure 7-19 HCM Connect both ISL ports

HCM represents ISL ports as shown in Figure 7-20 on page 213.




Figure 7-20 HCM ISL ports display

15. Double-click Switch #65 object to list all the defined ports.
16. Highlight Port 92 on Switch #65. Click **Connect**.
17. Navigate through the processors, CHPIDs, and Switches until you find the test processor.
18. Expand the CHPIDs, highlight CHPID 21, and then click **Connect** → **Physical Only**.
19. Perform the same action for Port A3 and CHPID 25. Refer to Figure 7-21 on page 214.




Figure 7-21 HCM Ports connected to Channel Paths

20. Click **OK**, and then Select **Edit** → **Controller**. Select the DASD controller for our test.

Presently our DASD subsystem has 8 FICON connections spread over 4 FICON directors plus 2 switched connections. For our example we will be adding 2 new cascaded FICON connections where the processor CHPIDs will connect through Switch ID #65 and the DS8000 will connect through Switch #66

21. Double-click the controller, and in the Interfaces box, highlight our two interfaces 0133 and 0333. Click **Connect**.

22. Highlight the first controller interface in the left panel (0133).

23. Locate Switch #66 in the right panel, and expand the port list until port 41 is located.

24. Highlight port 31, and click **Connect**.

25. Perform the same action for Controller interface 0333. Refer to Figure 7-22 on page 215.

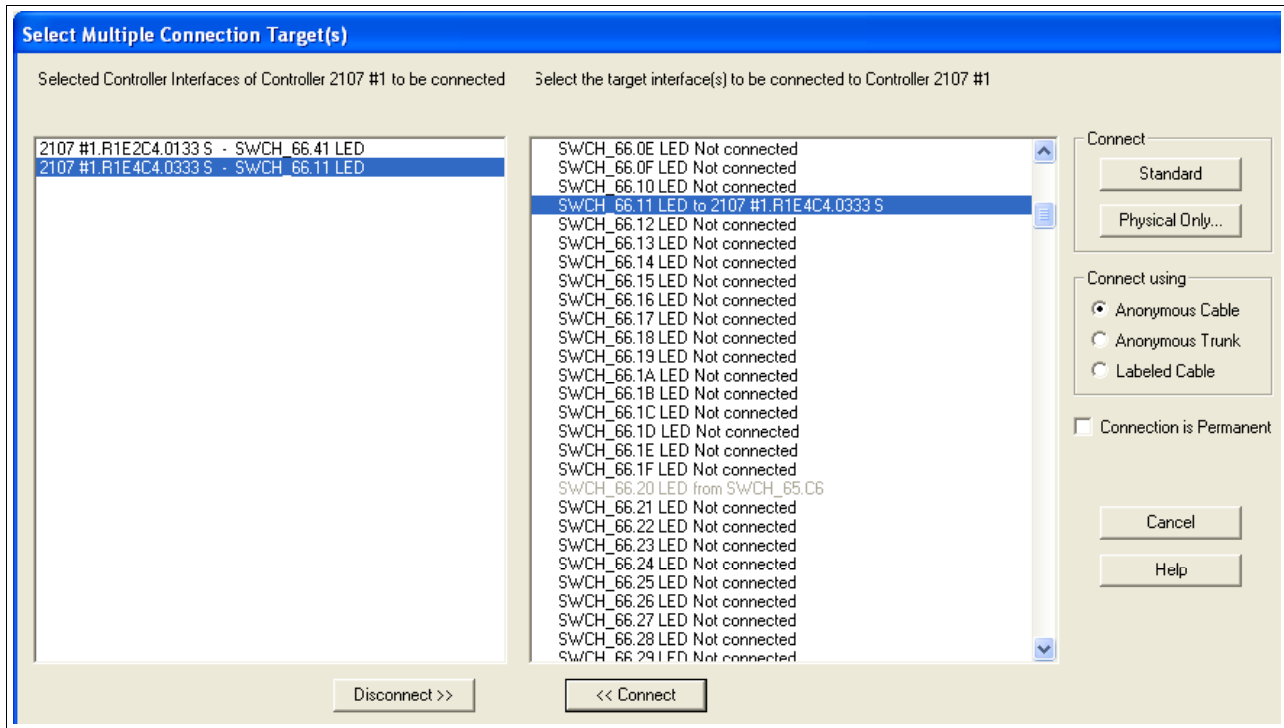


Figure 7-22 HCM Connecting to DS8000 interfaces to Switch Ports

26. Click **Standard** to continue to the Affected Control Units panel.
27. Highlight CU 9000, and then click **Connect**.
28. Highlight CHPID 21, and then click **Add**. See Figure 7-23.

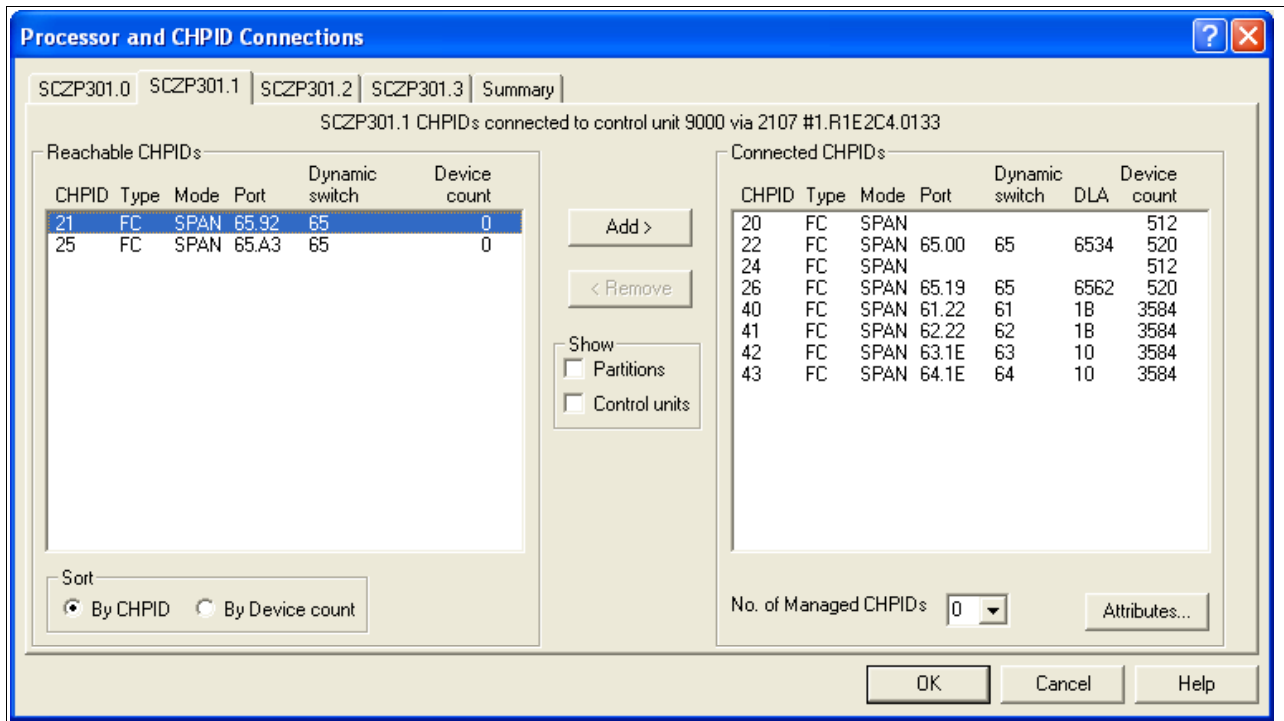


Figure 7-23 HCM Adding CHPIDs to CU definitions

29. Click **OK** to perform the same action for CU 9100.

30. Click **Close** to add CHPID 26 to CU 9000 and CU9100 through the interface 0233.

The final step looks like Figure 7-24. Observe the 2 byte Destination Link Addresses (DLA). Refer to Figure 7-24. The connections marked in red show:

- CHPID 21 -> Switch #65 / Port 92 -> Switch #66 / Port 41 -> CU interface 0133
- CHPID 25 -> Switch #65 / Port A3 -> Switch #66 / Port 11 -> CU interface 0333

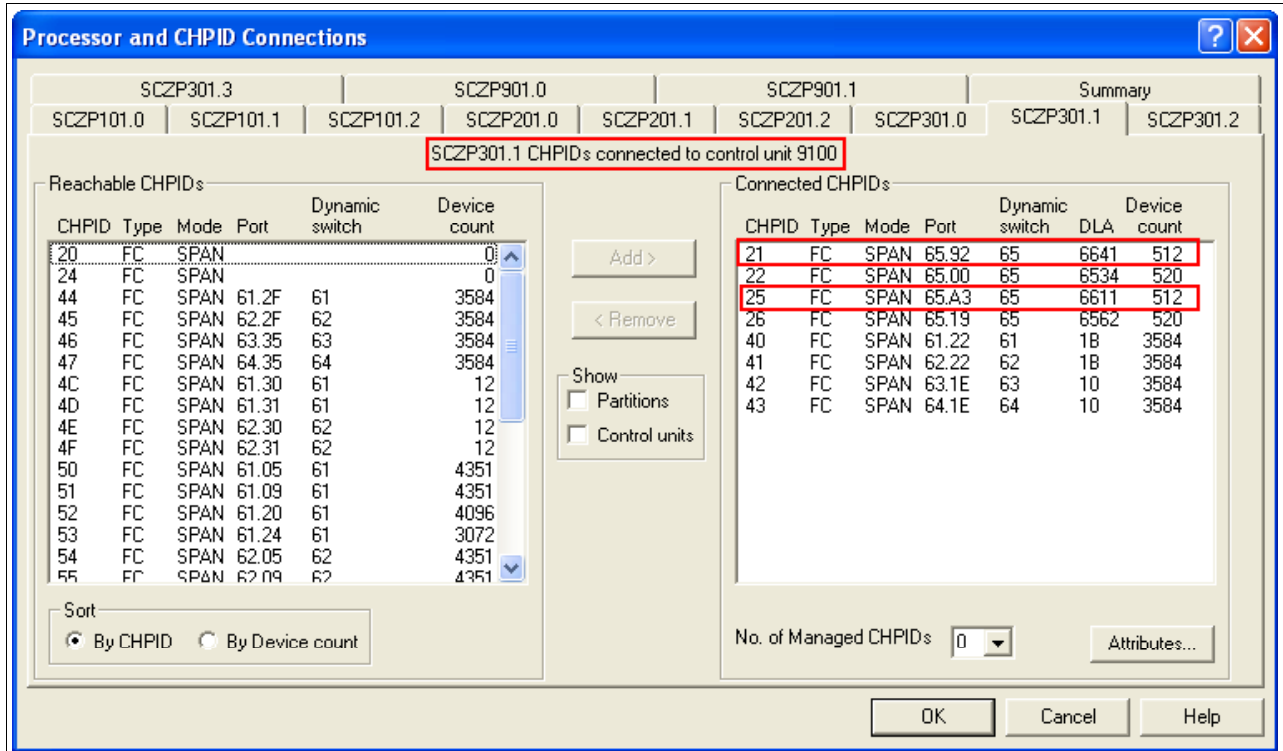


Figure 7-24 HCM Final step with adding CHPIDs to CU definitions

31. Click **OK** to return to the main HCM display and observe the new connections into interfaces 0133 and 0333. Refer to Figure 7-25 on page 217.



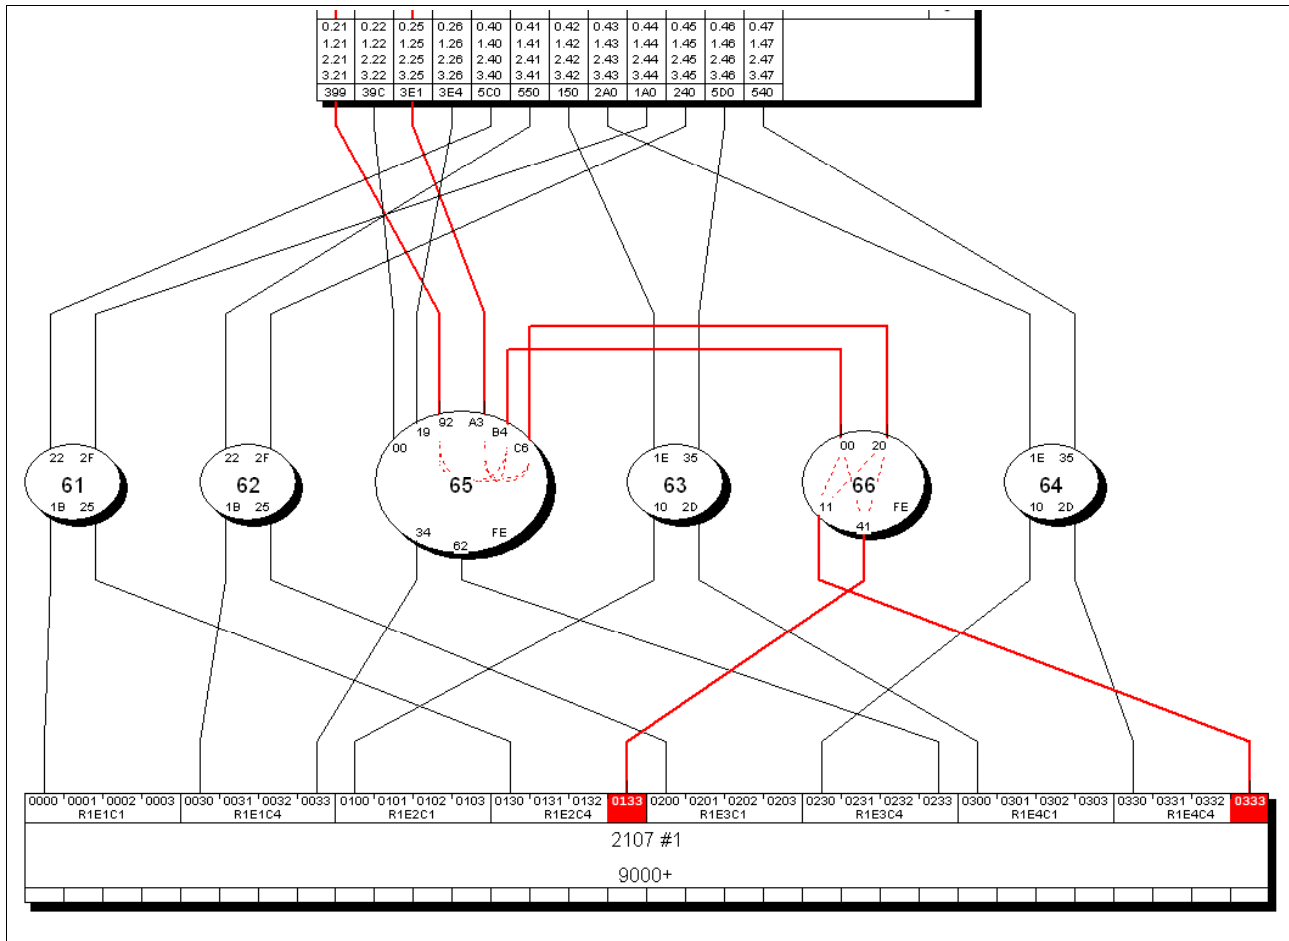


Figure 7-25 HCM DASD point-to-point connections

32. Ensure that the device definitions for the Operating System (software) are correct.
33. Double-click the controller box.
34. Click **Control Units**.
35. Highlight CU 9000, and then click **I/O Devices**.
36. Highlight all of the 3390B devices (9000-907F), and then click **OS Configs**.
37. Highlight TEST2094, and then click **Define**. Refer to Figure 7-26 on page 218.

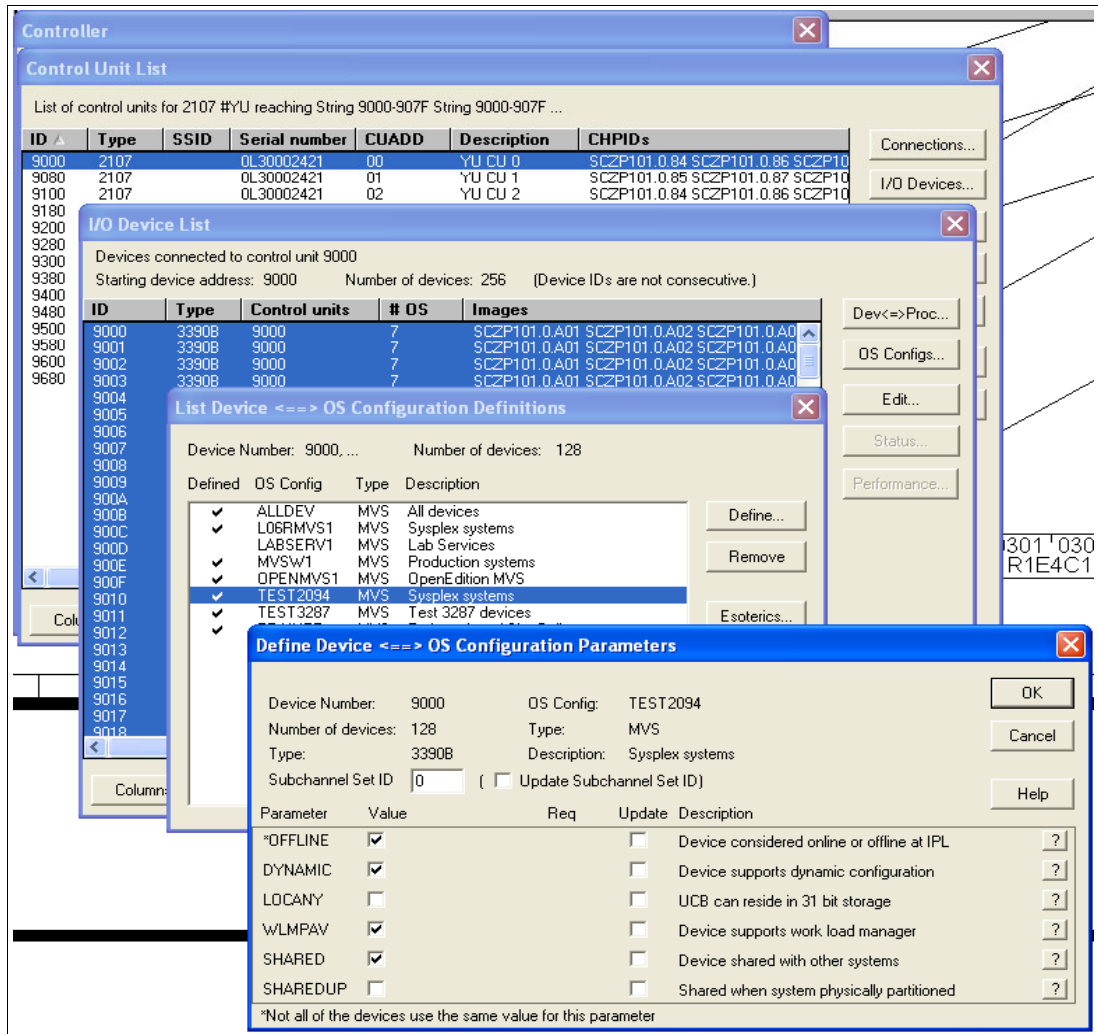


Figure 7-26 HCM OS Configuration Parameters

38. Verify that the Device Parameter values are correct.
39. Click **OK**, and then click **Esoterics**. Verify that the devices are defined to the correct Esoterics.
40. Perform the same operation for the 3390A devices for CU 9000 and for CU 9100 and Devices 9100-917F.

#### 7.4.4 How the definition looks as IOCP statements

The following examples show the CHPID, CNTLUNIT, and IODEVICE statements if all of the definition tasks based on the configuration shown in Figure 7-1 on page 198 are finished.

First, we explain how the CHPIDs are defined. Example 7-1 on page 219 displays the CHPID statement and its keywords.

*Example 7-1 CHPID definition for cascaded FICON configuration*

---

```
CHPID PATH=(CSS(0,1,2,3),21),SHARED,*
PARTITION=((CSS(3),(A31,A34,A35),(=))),*
NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),(=)),(CSS(1),*
(A1B,A1E,A1F),(=)),(CSS(2),(A2E,A2F),(=))),SWITCH=65,*
PCHID=399,TYPE=FC
CHPID PATH=(CSS(0,1,2,3),25),SHARED,*
PARTITION=((CSS(3),(A31,A34,A35),(=))),*
NOTPART=((CSS(0),(A08,A09,A0C,A0D,A0E,A0F),(=)),(CSS(1),*
(A1B,A1E,A1F),(=)),(CSS(2),(A2E,A2F),(=))),SWITCH=65,*
PCHID=3E1,TYPE=FC
```

---

In Example 7-1:

- ▶ The PATH keyword in the CHPID statement defines CHPID 21 and CHPID 25 in CSS(1) as SHARED.
- ▶ The NOTPART keyword specifies the LPARs which do not have access to CHPID 21 and CHPID 25
- ▶ The SWITCH keyword shows that this CHPID is connected to Switch #65
- ▶ The PCHID keyword specifies the physical channel identification number associated with CHPID 21 and CHPID 25

With the CHPIDs defined, we next show how the Control Units attached to the CHPIDs are defined. Example 7-2 displays the CNTLUNIT statement and keywords.

*Example 7-2 CNTLUNIT definition for a cascaded FICON configuration*

---

```
CNTLUNIT CUNUMBR=0065,PATH=((CSS(1),22,26)),*
UNITADD=((00,001)),LINK=((CSS(1),65FE,65FE)),UNIT=2032
CNTLUNIT CUNUMBR=0066,PATH=((CSS(1),22,26)),*
UNITADD=((00,001)),LINK=((CSS(1),66FE,66FE)),UNIT=2032
CNTLUNIT CUNUMBR=9000,*
PATH=((CSS(0),40,41,42,43),(CSS(1),40,41,42,43,22,26,21,*
25),(CSS(2),40,41,42,43),(CSS(3),40,41,42,43)),*
UNITADD=((00,256)),*
LINK=((CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,6534,6562*
,6641,6611),(CSS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)),*
CUADD=0,UNIT=2107
CNTLUNIT CUNUMBR=9100,*
PATH=((CSS(0),40,41,42,43),(CSS(1),40,41,42,43,22,26,21,*
25),(CSS(2),40,41,42,43),(CSS(3),40,41,42,43)),*
UNITADD=((00,256)),*
LINK=((CSS(0),1B,1B,10,10),(CSS(1),1B,1B,10,10,6534,6562*
,6641,6611),(CSS(2),1B,1B,10,10),(CSS(3),1B,1B,10,10)),*
CUADD=2,UNIT=2107
```

---

Included in this display are both the Switch CUP control unit definitions along with the DS8000 control unit definitions:

- ▶ The CUNUMBR keyword shows definitions for 0065, 0066, 9000, and 9100.
- ▶ The PATH keyword shows CSS(1) as having definitions for CHPID 22 and CHPID 26 for the switch CU definitions and CHPIDs 21 and 25 for the DS8000 CU definitions.
- ▶ The LINK keyword shows the 2 byte switch and destination link address (DLA) that the control units are defined to. Observe the FE (or CUP port) for the Switch definitions.
- ▶ The CUADD keyword (representing the logical control unit or LCU) is set to 0 for CU 9000 and 2 for 9100. By pressing F20 in the panel shown in Figure 7-8 on page 206, you can see the CUADD value set.

After the CHPIDs and the CUs are defined, the next step is to define the devices owned by the CUs. Example 7-3 displays the IODEVICE statement and keywords.

*Example 7-3 IODEVICE statement for cascaded FICON configuration*

---

```

IODEVICE ADDRESS=065,UNITADD=00,CUNUMBR=(0065),STADET=Y,      *
          UNIT=2032
IODEVICE ADDRESS=066,UNITADD=00,CUNUMBR=(0066),STADET=Y,      *
          UNIT=2032
IODEVICE ADDRESS=(9000,128),CUNUMBR=(9000),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9000,128),UNITADD=80,CUNUMBR=(9000),        *
          STADET=Y,SCHSET=1,UNIT=3390A
IODEVICE ADDRESS=(9100,128),CUNUMBR=(9100),STADET=Y,UNIT=3390B
IODEVICE ADDRESS=(9100,128),UNITADD=80,CUNUMBR=(9100),        *
          STADET=Y,SCHSET=1,UNIT=3390A

```

---

In Example 7-3:

- ▶ The ADDRESS keyword shows device address 065 or 066 for the Switch CUP devices along with the DS8000 devices 9000 for 128 (9000-907F) or 9100 for 128 (9100-917F)
- ▶ The CUNUMBR keyword associates the IODEVICE statement with the Control Unit number 0065, 0066 or 9000, 9100
- ▶ The UNIT keyword shows the device type of 3390B (for the base devices) or 3390A (for the HyperPAV alias devices)
- ▶ The UNITADD keyword is used when the Unit Address (UA) does not start at 00
- ▶ The SCHSET keyword shows that these devices are defined in Subchannel Set 1 (SS1)

**Note:** For CU 9000 and CU 9100, the device allocation is defined in the following way:

- ▶ 3390B devices 9000-907F, CUADD=0, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9000-907F, CUADD=0, UA = 80-FF, SCHSET = 1
- ▶ 3390B devices 9100-917F, CUADD=2, UA = 00-7F, SCHSET = 0
- ▶ 3390A devices 9100-917F, CUADD=2, UA = 80-FF, SCHSET = 1

## 7.4.5 Defining a FICON Channel-to-Channel (FCTC) connection

In addition to storage control units, storage devices, and the CUPs in the FICON Directors, we defined some cascaded FICON Channel-to-Channel (FCTC) CUs and devices to allow communication between LPARs on the same *and* different processors. For more information about considerations regarding FCTCs, refer to Appendix G, “Adding FICON CTC connections” on page 489.

Based on those considerations, we must configure FCTCs in a cascaded FICON environment. Figure 7-27 on page 221 illustrates our configuration. To simplify the illustration, four CUs and two LPARs are shown in this sample configuration, although more than four control units can be defined.




Figure 7-27 FCTC configuration (cascaded)

Both z196 servers are attached to cascaded FICON Directors:

- ▶ From processor SCZP301:
  - CHPID 21 to port 92 on switch 65 and CHPID 25 to port A3 on switch 65
- ▶ From processor SCZP302:
  - CHPID 23 to port 30 on switch 66 and CHPID 27 to port 50 on switch 66
- ▶ Switch 65 and switch 66 are connected using Inter-Switch Links (ISLs):
  - Port B4 from switch 65 connects to port 00 on switch 66
  - Port C6 from switch 65 connects to port 20 on switch 66
- ▶ All other LPARs communicate with:
  - LPAR A11 on SCZP301 using CU 4114 on CHPID 21 (pathin)
  - LPAR A13 on SCZP301 using CU 4134 on CHPID 21 (pathin)
  - LPAR A01 on SCZP302 using CU 4014 on CHPID 23 (pathin)
  - LPAR A03 on SCZP302 using CU 4034 on CHPID 23 (pathin)
- ▶ Outbound communications for:
  - LPAR A11 on SCZP301 communicates with all other LPARs using CU 5114 on CHPID 25 (pathout)
  - LPAR A13 on SCZP301 communicates with all other LPARs using CU 5134 on CHPID 25 (pathout)
  - LPAR A01 on SCZP302 communicates with all other LPARs using CU 5014 on CHPID 27 (pathout)
  - LPAR A03 on SCZP302 communicates with all other LPARs using CU 5034 on CHPID 27 (pathout)

Here is the CTC definition panel in HCM for the FCTC connection between SCZP301 CHPID 21 and CHPID 25. Refer to Figure 7-28 on page 222.

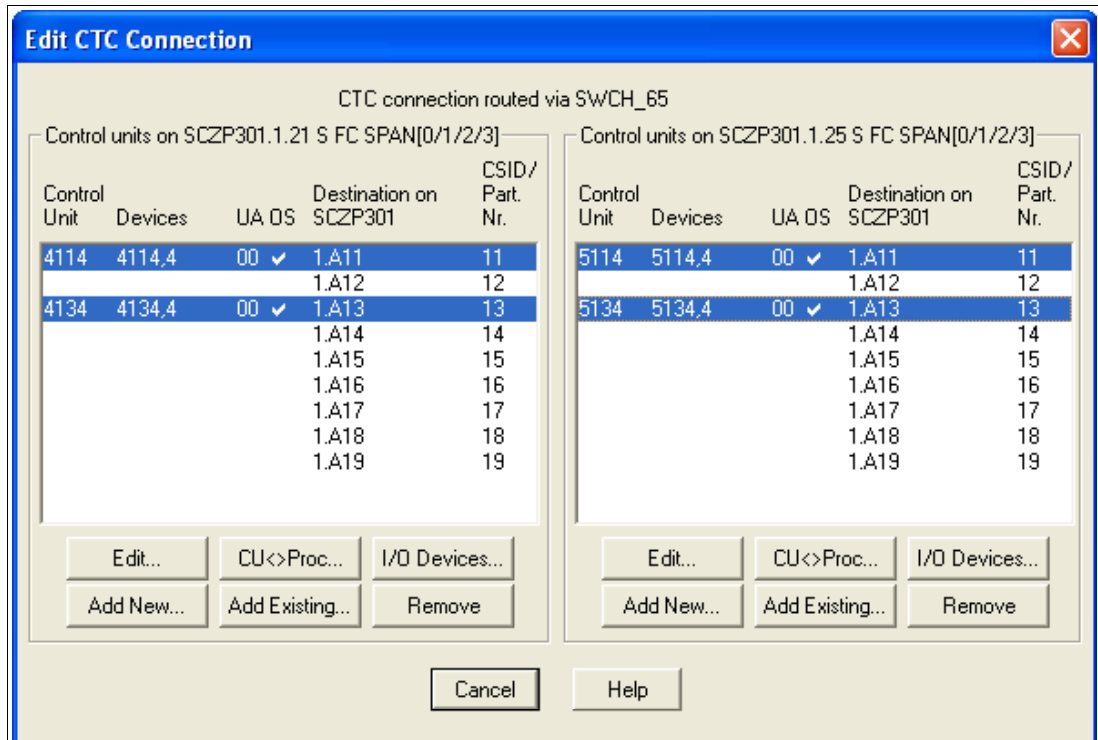


Figure 7-28 HCM FCTC definitions between CHPID 21 and CHPID 25 (switched)

Here is the CTC definition panel in HCM for the FCTC connection between SCZP302 CHPID 23 and CHPID 27. Refer to Figure 7-29.

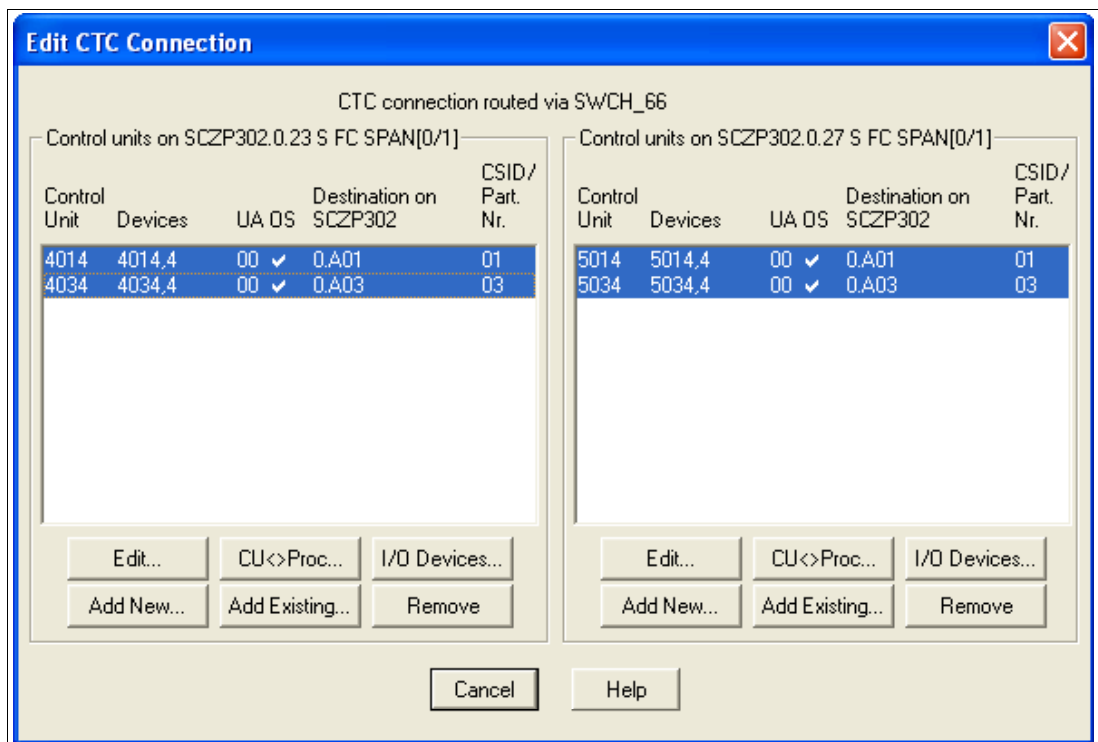


Figure 7-29 HCM FCTC definitions between CHPID 23 and CHPID 27 (switched)

If we display the CTC definition panel in HCM for the FCTC connection between SCZP301 CHPID 21 and SCZP302 CHPID 27, notice that the panel indicates that the connection is through cascaded switches. Refer to Figure 7-30.

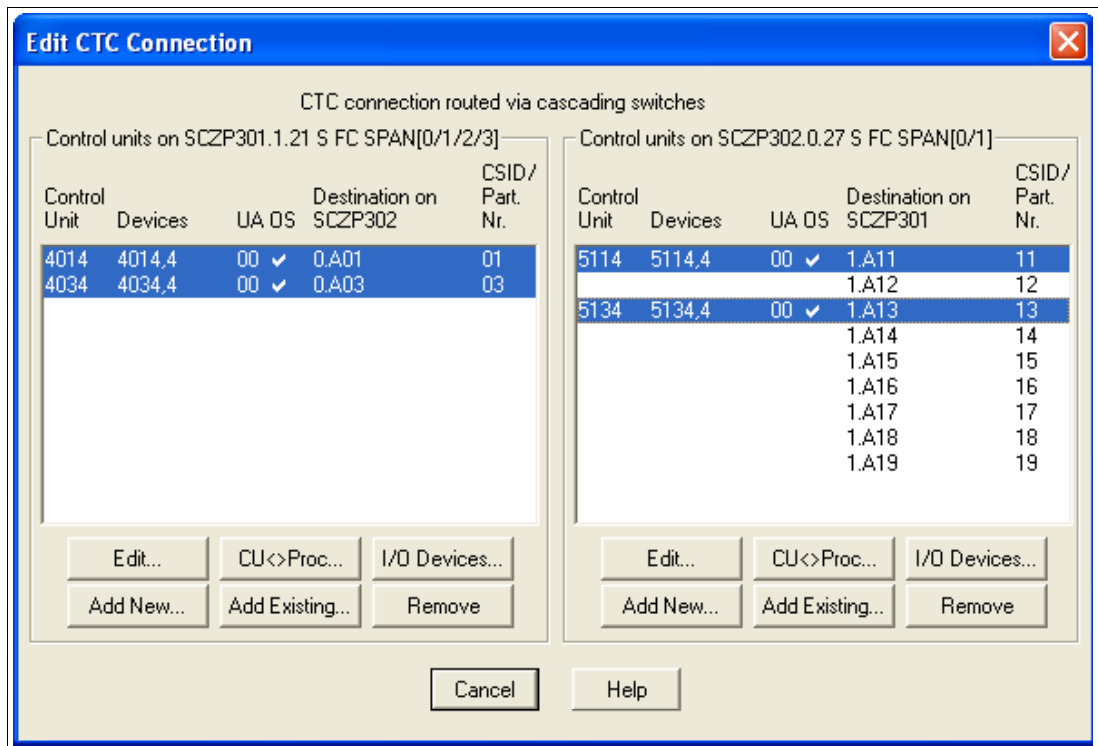


Figure 7-30 HCM FCTC definitions between CHPID 21 and CHPID 27 (cascaded)

Here is a how HCM displays switched/cascaded FCTC definitions graphically. Refer to Figure 7-31 on page 224.




Figure 7-31 HCM FCTC switches and cascaded connections

In Figure 7-31, we can see in the highlighted connection that:

- SCZP301 CHPID 21 has a link to SCZP301 CHPID 25 (switched)
- SCZP301 CHPID 21 has a link to SCZP302 CHPID 27 through either ISL connection B4 <-> 00 or C6 <-> 20 (cascaded)

Example 7-5 on page 225 shows the IOCP statements for these FCTC control units and device definitions defined on processor SCZP301.

*Example 7-4 FCTC configuration for SCZP301*

```

CNTLUNIT CUNUMBR=4014,PATH=((CSS(1),21)),UNITADD=((00,004)), *
        LINK=((CSS(1),6650)),CUADD=1,UNIT=FCTC
IODEVICE ADDRESS=(4014,004),UNITADD=00,CUNUMBR=(4014), *
```



STADET=Y,UNIT=FCTC

```
CNTLUNIT CUNUMBR=4034,PATH=((CSS(1),21)),UNITADD=((00,004)), *
LINK=((CSS(1),6650)),CUADD=3,UNIT=FCTC
IODEVICE ADDRESS=(4034,004),UNITADD=00,CUNUMBR=(4034), *
STADET=Y,UNIT=FCTC

CNTLUNIT CUNUMBR=4114,PATH=((CSS(1),21)),UNITADD=((00,004)), *
LINK=((CSS(1),65A3)),CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(4114,004),UNITADD=00,CUNUMBR=(4114), *
STADET=Y,NOTPART=((CSS(1),A11,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=4134,PATH=((CSS(1),21)),UNITADD=((00,004)), *
LINK=((CSS(1),65A3)),CUADD=13,UNIT=FCTC
IODEVICE ADDRESS=(4134,004),UNITADD=00,CUNUMBR=(4134), *
STADET=Y,NOTPART=((CSS(1),A13,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=5014,PATH=((CSS(1),25)),UNITADD=((00,004)), *
LINK=((CSS(1),6630)),CUADD=1,UNIT=FCTC
IODEVICE ADDRESS=(5014,004),UNITADD=00,CUNUMBR=(5014), *
STADET=Y,UNIT=FCTC

CNTLUNIT CUNUMBR=5034,PATH=((CSS(1),25)),UNITADD=((00,004)), *
LINK=((CSS(1),6630)),CUADD=3,UNIT=FCTC
IODEVICE ADDRESS=(5034,004),UNITADD=00,CUNUMBR=(5034), *
STADET=Y,UNIT=FCTC

CNTLUNIT CUNUMBR=5114,PATH=((CSS(1),25)),UNITADD=((00,004)), *
LINK=((CSS(1),6592)),CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(5114,004),UNITADD=00,CUNUMBR=(5114), *
STADET=Y,NOTPART=((CSS(1),A11,A1B,A1E,A1F)),UNIT=FCTC

CNTLUNIT CUNUMBR=5134,PATH=((CSS(1),25)),UNITADD=((00,004)), *
LINK=((CSS(1),6592)),CUADD=13,UNIT=FCTC
IODEVICE ADDRESS=(5134,004),UNITADD=00,CUNUMBR=(5134), *
STADET=Y,NOTPART=((CSS(1),A13,A1B,A1E,A1F)),UNIT=FCTC
```

---

Example 7-5 shows the IOCP statements for these FCTC control units and device definitions defined on processor SCZP302.

*Example 7-5 FCTC configuration for SCZP302*

---

```
CNTLUNIT CUNUMBR=4014,PATH=((CSS(0),23)),UNITADD=((00,004)), *
LINK=((CSS(0),6650)),CUADD=1,UNIT=FCTC
IODEVICE ADDRESS=(4014,004),UNITADD=00,CUNUMBR=(4014), *
STADET=Y,PARTITION=((CSS(0),A03)),UNIT=FCTC

CNTLUNIT CUNUMBR=4034,PATH=((CSS(0),23)),UNITADD=((00,004)), *
LINK=((CSS(0),6650)),CUADD=3,UNIT=FCTC
IODEVICE ADDRESS=(4034,004),UNITADD=00,CUNUMBR=(4034), *
STADET=Y,PARTITION=((CSS(0),A01)),UNIT=FCTC

CNTLUNIT CUNUMBR=4114,PATH=((CSS(0),23)),UNITADD=((00,004)), *
LINK=((CSS(0),65A3)),CUADD=11,UNIT=FCTC
IODEVICE ADDRESS=(4114,004),UNITADD=00,CUNUMBR=(4114), *
```

```

          STADET=Y, PARTITION=((CSS(0),A03)), UNIT=FCTC

CNTLUNIT CUNUMBR=4134, PATH=((CSS(0),23)), UNITADD=((00,004)), *
          LINK=((CSS(0),65A3)), CUADD=13, UNIT=FCTC
IODEVICE ADDRESS=(4134,004), UNITADD=00, CUNUMBR=(4134), *
          STADET=Y, PARTITION=((CSS(0),A01)), UNIT=FCTC

CNTLUNIT CUNUMBR=5014, PATH=((CSS(0),27)), UNITADD=((00,004)), *
          LINK=((CSS(0),6630)), CUADD=1, UNIT=FCTC
IODEVICE ADDRESS=(5014,004), UNITADD=00, CUNUMBR=(5014), *
          STADET=Y, PARTITION=((CSS(0),A03)), UNIT=FCTC

CNTLUNIT CUNUMBR=5034, PATH=((CSS(0),27)), UNITADD=((00,004)), *
          LINK=((CSS(0),6630)), CUADD=3, UNIT=FCTC
IODEVICE ADDRESS=(5034,004), UNITADD=00, CUNUMBR=(5034), *
          STADET=Y, PARTITION=((CSS(0),A01)), UNIT=FCTC

CNTLUNIT CUNUMBR=5114, PATH=((CSS(0),27)), UNITADD=((00,004)), *
          LINK=((CSS(0),6592)), CUADD=11, UNIT=FCTC
IODEVICE ADDRESS=(5114,004), UNITADD=00, CUNUMBR=(5114), *
          STADET=Y, PARTITION=((CSS(0),A03)), UNIT=FCTC

CNTLUNIT CUNUMBR=5134, PATH=((CSS(0),27)), UNITADD=((00,004)), *
          LINK=((CSS(0),6592)), CUADD=13, UNIT=FCTC
IODEVICE ADDRESS=(5134,004), UNITADD=00, CUNUMBR=(5134), *
          STADET=Y, PARTITION=((CSS(0),A01)), UNIT=FCTC

```

---

In Example 7-5 on page 225:

- ▶ CUNUMBR keyword shows CU 4114 and 5114 for LPAR A11 and CU 4134 and 5134 for LPAR A13
- ▶ CUNUMBR keyword shows CU 4014 and 5014 for LPAR A01 and CU 4034 and 5034 for LPAR A03
- ▶ PATH keyword shows the CSS and CHPID that the CU is defined to
- ▶ LINK keyword show the destination link address (DLA) where the CU can be located
- ▶ CUADD keyword shows the logical control unit number (LCU), which must equal the LPAR ID
- ▶ NOTPART keyword shows the partitions on this processor which cannot access these devices. This is controlled using the Explicit Device Candidate List (EDCL) option in the Operating system configuration definition.
- ▶ PARTITION keyword shows the partitions on this processor which have access to these devices. This is controlled using the Explicit Device Candidate List (EDCL) option in the Operating system configuration definition. The rules we followed for device numbering are described in “FCTC device numbering scheme” on page 490.

Figure 7-32 on page 227 illustrates the logical view of the FCTC configuration and the data path between two corresponding FCTC devices. To simplify the illustration, only two CUs per processor are shown.

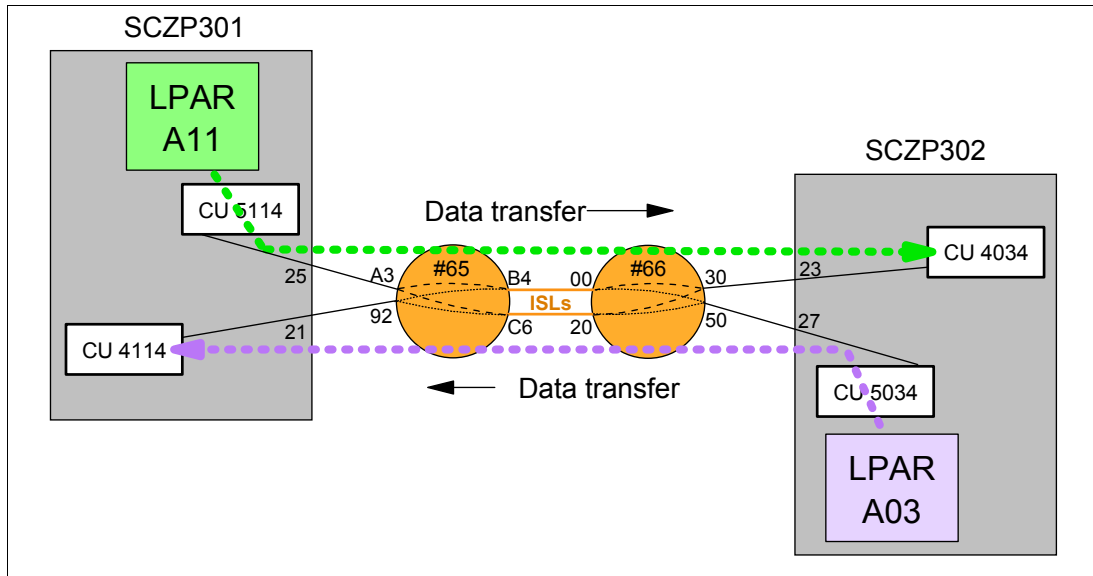


Figure 7-32 FCTC data transfer (cascaded)

Data is transferred over the FICON link in both directions between the CUs and the logical partitions (LPARs):

- ▶ LPAR A11 sends data to LPAR A03 through CU 4034.
- ▶ LPAR A11 receives data from LPAR A03 through CU 4114.

The reverse applies to LPAR A03 when data is sent to or received from LPAR A11.

## 7.5 Building and activating the IODF

After the definition of the desired configuration is finished, a production IODF needs to be built and saved to the IODF volume and the IODF data set and written to an IOCDS on the Support Element. Up to four IOCDS files can be stored in a zEnterprise 196 server. This allows us to plan and define future I/O configurations and store them on the support element.

The tasks that are required to build a production IODF, save I/O definition data, and activate changes dynamically are performed through HCD. Refer to *HCD User's Guide*, SC33-7988, for detailed descriptions of all activation procedures.

### 7.5.1 Building the production IODF

To build the production IODF:

1. Select the current work IODF using HCD. We use SYS6.IODF07.WORK for our example.
2. From the HCD main menu, select **Option 2.1 Build production I/O definition file**.
3. Review the Message List for any messages that require attention.
4. Press PF3 to continue to the Build Production I/O Definition File panel.
5. Type the Production IODF name in the Production IODF name field. Refer to Figure 7-33 on page 228.

```

----- Build Production I/O Definition File -----

Specify the following values, and choose how to continue.

Work IODF name . . . : 'SYS6.IODF07.WORK'

Production IODF name . 'SYS6.IODF07' _____
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 7-33 HCD Build Production IODF

6. Press Enter and confirm the Descriptor field names. See Figure 7-34.

```

----- Define Descriptor Fields -----

Specify or revise the following values.

Production IODF name . : 'SYS6.IODF07'

Descriptor field 1 . . . SYS6
Descriptor field 2 . . . IODF07

```

Figure 7-34 HCD Descriptor Fields

7. Press Enter to complete the build process. The message Production IODF SYS6.IODF07 created confirms this.
8. Return to the HCD main menu, and select **Option 2.11 Build and manage S/390 microprocessor IOCDs and IPL attributes**.
9. A list of processors that are accessible from the system that HCD is running on is displayed. Refer to Figure 7-35. Type s next to the processor in our example.

```

                                S/390 Microprocessor Cluster List
Command ==> _____ Row 1 of 6
                                Scroll ==> CSR

Select one or more CPCs, then press Enter.

-----CPC----- IODF
/ SNA Address      Type  Model Processor ID
# IBM390PS.H40
# IBM390PS.TC8M
_ USIBMSC.SCZP101  2094 S18   SCZP101
_ USIBMSC.SCZP201  2097 E26   SCZP201
s USIBMSC.SCZP301  2817 M32   SCZP301
_ USIBMSC.SCZP901  2084 C24   SCZP901

```

Figure 7-35 HCD Writing an IOCDs from an IODF

The next panel shows a summary of all IOCDSs for this processor. Here we can see the current IOCDS that is in use by this processor indicated by the YES-POR under the Write Protect column. Refer to Figure 7-36.

10.Type u next to IOCDS A0 to update this IOCDS with our new IODF information.

```

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

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

/ IOCDS      Name      Type      Status      -----Token Match----- Write
u A0.SCZP301  IODF77  LPAR      Alternate  No          No          No
_ A1.SCZP301  IODF00  LPAR      Alternate  No          No          No
_ A2.SCZP301  IODF05  LPAR      Alternate  No          No          No
_ A3.SCZP301  IODF06  LPAR      POR        Yes         No          Yes-POR

```

Figure 7-36 HCD IOCDS List

11.Press enter to proceed to the next panel. Refer to Figure 7-37.

12.Type IODF07 in the Title1 field. Leave the other options as defaults.

```

----- Build IOCDSs -----
Row 1 of 1
Command ===> _____ Scroll ===> CSR

Specify or revise the following values.

IODF name . . . . . : 'SYS6.IODF07'

Title1 . IODF07 _____
Title2 : SYS6.IODF07 - 2011-11-21 15:01

IOCDS      Switch IOCDS      Write IOCDS in
A0.SCZP301 No          preparation of upgrade
No

```

Figure 7-37 HCD Build IOCDS

13.Press enter to continue to the final panel. See Figure 7-38.

```

----- Job Statement Information -----

Specify or revise the job statement information.

Job statement information
//IOCDSA0 JOB (ACCOUNT),'NAME',REGION=128M
//*
//*
//*
//*

```

Figure 7-38 HCD IOCDS Job Statement panel

14. Ensure that the JCL statements meet your site's requirements, and press Enter to submit the Write IOCDS job.
15. After the job completes, verify the system output and look for the message in Example 7-6.

*Example 7-6 HCD IOCDS written successfully*

---

CBDA674I IOCP successfully completed for A0.SCZP301.

---

16. Return to HCD Option 2.11, and navigate back to the IOCDS List for this processor. See Figure 7-39.  
Observe that IOCDS A0 now has a name of IODF07 and that there is now a Token Match of Yes under the IOCDS/Proc column.

| IOCDS List   |               |             |                  | Row 1 of 4 More: >          |
|--|---------------|-------------|------------------|-----------------------------|
| Command ==> _____                                  |               |             |                  | Scroll ==> CSR              |
| Select one or a group of IOCDSs, then press Enter. |               |             |                  |                             |
| / IOCDS  | Name          | Type        | Status           | -----Token Match----- Write |
| - <b>A0.SCZP301</b>                                | <b>IODF07</b> | <b>LPAR</b> | <b>Alternate</b> | <b>No Yes No</b>            |
| - A1.SCZP301                                       | IODF00        | LPAR        | Alternate        | No No No                    |
| - A2.SCZP301                                       | IODF05        | LPAR        | Alternate        | No No No                    |
| - A3.SCZP301                                       | IODF06        | LPAR        | POR              | Yes No Yes-POR              |

Figure 7-39 HCD IOCDS List after IODF update

## 7.5.2 Activating the production IODF

There are two ways to activate an IODF. One way is to use the HCD panels in Option 2 and another way is using the ACTIVATE system command. We demonstrate using the ACTIVATE system command:

1. Access a console or SDSF under TSO for this system where commands are authorized to be entered.
2. Enter **D IOS,CONFIG** to see what IODF the hardware and software use for this system. Refer to Figure 7-40:
  - Active IODF data set: Shows which IODF the Software (or LPAR) is currently using
  - Source: Shows which IODF the Hardware (or HSA) for processor SCZP301 is currently using

```

-D IOS,CONFIG
IOS506I 15.43.13 I/O CONFIG DATA 286
ACTIVE IODF DATA SET = SYS6.IODF06
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-20 11:41:25 SYS6      IODF06
ACTIVE CSS: 1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 7-40 MVS IODF system display

3. Enter **ACTIVATE IODF=07,TEST** to test the new IODF and verify that the changes we made to the IODF are reflected accurately in the test display. Refer to Figure 7-41.

This command tests changes being made to the Hardware or HSA.

```
-ACTIVATE IODF=07,TEST  
IOS500I ACTIVATE RESULTS 300  
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE  
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:  
      1.0065-1.0066  
      COMPID=SC1XL  
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR  
      SCZP301:  
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917  
      COMPID=SC1XL  
NOTE = A888,F  
      COMPID=SC1XL  
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR  
      SCZP301: 1.0065-1.0066  
      COMPID=SC1XL  
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR  
      SCZP301: 1.9000,1.9100  
      COMPID=SC1XL  
NOTE = A879,FOLLOWING CHANNEL PATHS ARE TO BE DELETED FROM PROCESSOR  
      SCZP301: 1.21,1.25,2.21,2.25,3.21,3.25,0.21,0.25  
      COMPID=SC1XL  
NOTE = A881,FOLLOWING CHANNEL PATHS ARE TO BE ADDED TO PROCESSOR  
      SCZP301: 0.21,0.25,1.21,1.25,2.21,2.25,3.21,3.25  
      COMPID=SC1XL  
NOTE = 0112,REQUEST CONTAINS DELETE(S), SPECIFY FORCE ON H/W ACTIVATE  
COMPID=SC1C3
```

Figure 7-41 MVS Test ACTIVATE display

4. Enter **ACTIVATE,IODF=07,SOFT,TEST** to test changes being made to the software or operating system (TEST2094) for this system. Refer to Figure 7-42.

```
-ACTIVATE IODF=07,SOFT,TEST  
IOS500I ACTIVATE RESULTS 329  
TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE  
NOTE = 0100,SOFTWARE-ONLY CHANGE  
      COMPID=SC1C3  
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS WILL BE OUT OF SYNC  
      COMPID=SC1C3
```

Figure 7-42 MVS Test ACTIVATE SOFT command

The responses to the display look positive as per the message **TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE**.

We are now ready to activate the new IODF to the hardware and software for this system.

Because our system (SC30) is part of a sysplex (COMMPLEX), we must first activate the IODF on those systems. The systems defined in this sysplex are SC30, SC31, SC32, and SC33. Refer to Figure 7-3 on page 203.

```

-D XCF
IXC334I 10.18.08 DISPLAY XCF 305
  SYSPLEX COMMPLEX:  SC30          SC31          SC32
                    SC33

```

Figure 7-43 MVS D XCF command

A SOFT ACTIVATE was performed already on systems SC32 and SC33.

- For system SC31, enter **ACTIVATE IODF=07,SOFT**.

```

-RO SC31,ACTIVATE IODF=07,SOFT
IOS1500I ACTIVATE RESULTS 718
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 0100,SOFTWARE-ONLY CHANGE
  COMPID=SC1C3
NOTE = 010B,H/W AND S/W CONFIGURATION DEFINITIONS ARE NOW OUT OF SYNC
  COMPID=SC1C3 COMPID=SC1C3

```

Figure 7-44 MVS ACTIVATE SOFT command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 7-44.

- Enter D IOS,CONFIG to verify that the activate for the LPAR SC31 completed.

```

-RO SC31,D IOS,CONFIG
IOS506I 15.52.46 I/O CONFIG DATA 720
ACTIVE IODF DATA SET = SYS6.IODF07
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-20 11:41:25 SYS6   IODF06
ACTIVE CSS: 1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 7-45 MVS D IOS.CONFIG command

Observe that the software is on SYS6.IODF07, and the hardware is on SYS6.IODF06. Refer to Figure 7-45.

- For system SC30, enter **ACTIVATE IODF=07**.



```

-RO SC30,ACTIVATE IODF=07,FORCE
IOS1500I ACTIVATE RESULTS 388
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = A887,FOLLOWING DEVICES ARE TO BE ADDED TO PROCESSOR SCZP301:
      1.0065-1.0066
      COMPID=SC1XL
NOTE = A888,FOLLOWING DEVICES ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301:
      1.9000-1.907F,1.19000-1.1907F,1.9100-1.917F,1.19100-1.1917
      COMPID=SC1XL
NOTE = A888,F
      COMPID=SC1XL
NOTE = A884,FOLLOWING CONTROL UNITS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 1.0065-1.0066
      COMPID=SC1XL
NOTE = A889,FOLLOWING CONTROL UNITS ARE TO BE MODIFIED FOR PROCESSOR
      SCZP301: 1.9000,1.9100
      COMPID=SC1XL
NOTE = A879,FOLLOWING CHANNEL PATHS ARE TO BE DELETED FROM PROCESSOR
      SCZP301: 1.21,1.25,2.21,2.25,3.21,3.25,0.21,0.25
      COMPID=SC1XL
NOTE = A881,FOLLOWING CHANNEL PATHS ARE TO BE ADDED TO PROCESSOR
      SCZP301: 0.21,0.25,1.21,1.25,2.21,2.25,3.21,3.25
      COMPID=SC1XL

```

Figure 7-46 MVS ACTIVATE IODF command

Observe the response **ACTIVATE COMPLETED SUCCESSFULLY**. Refer to Figure 7-46.

**Note:** FORCE was specified in this ACTIVATE command because the CHPIDs in this IODF were previously defined to this processor (SCZP301), but we needed to add Switch information to them, which required the Activate process to take the Channels offline then online again.

8. Enter D IOS,CONFIG to verify that the activate for the LPAR SC30 completed.

```

-RO SC30,D IOS,CONFIG
IOS506I 15.56.07 I/O CONFIG DATA 406
ACTIVE IODF DATA SET = SYS6.IODF07
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE      TIME      DESCRIPTION
SOURCE: SCZP301 11-11-21 15:01:54 SYS6   IODF07
ACTIVE CSS:  1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE

```

Figure 7-47 MVS D IOS,CONFIG command

Observe that the software is on SYS6.IODF07, and the hardware is now on SYS6.IODF07. Refer to Figure 7-47.

The IOCDS must be switched to IOCDS A0 so that the Reset Profile for this processor SCZP301 points to the correct IODF information should there be a scheduled or unscheduled POR of this processor.

9. Enter **ACTIVATE ACTIOCDs=A0** to switch the IOCDs.

```
-RO SC30,ACTIVATE ACTIOCDs=A0
IOS500I ACTIVATE RESULTS 415
ACTIVATE COMPLETED SUCCESSFULLY
NOTE = 010C,THE ACTIVE IOCDs HAS BEEN SWITCHED TO A0
COMPID=SC1C3
```

Figure 7-48 MVS ACTIVATE ACTIOCDs command

Observe the message THE ACTIVE IOCDs HAS NOW BEEN SWITCHED TO A0. Refer to Figure 7-48.

10.If we return to HCD and select the following items, we see that IOCDs A0 is now selected and that a 'Yes' response is displayed for all columns. Refer to Figure 7-49:

- SYS6.IODF07
- Option 2.11
- our test processor SCZP301

```
IOCDs List Row 1 of 4 More: >
Command ==> _____ Scroll ==> CSR

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

-----Token Match----- Write
/ IOCDs Name Type Status IOCDs/HSA IOCDs/Proc. Protect
_ A0.SCZP301 IODF07 LPAR POR Yes Yes Yes-POR
_ A1.SCZP301 IODF00 LPAR Alternate No No No
_ A2.SCZP301 IODF05 LPAR Alternate No No No
_ A3.SCZP301 IODF06 LPAR Alternate No No No
```

Figure 7-49 HCD IOCDs List after new IODF activate

By pressing F20, more information can be displayed in the IOCDs List panel. See Figure 7-50.

```
IOCDs List Row 1 of 4 More: <
Command ==> _____ Scroll ==> CSR

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

--Last Update--
/ IOCDs Date Time IOCDs Configuration Token Information
_ A0.SCZP301 2011-11-21 15:39 SCZP301 15:01:54 11-11-21 SYS6 IODF07
_ A1.SCZP301 2011-11-16 17:53 SCZP301 09:25:30 11-11-10 SYS6 IODF00
_ A2.SCZP301 2011-11-16 18:12 SCZP301 16:39:06 11-11-16 SYS6 IODF05
_ A3.SCZP301 2011-11-20 12:09 SCZP301 11:41:25 11-11-20 SYS6 IODF06
```

Figure 7-50 HCD IOCDs List after pressing F20

Now that the hardware component (HSA) of the IODF is synchronized, we must make sure that the software component is also synchronized.

The volume where the IODF was saved before (IODFPK) also contains a dataset called SYS#.IPLPARM. This dataset contains LOADxx members, which are used at IPL startup

to determine certain system parameters. It also tells the system being IPLed, which IODF to use.

When an LPAR is IPLed, a Load (IPL) address and a Load parameter (loadparm) are specified

The Load Profile looks like Figure 7-51 for this system.

|   |  |  |
|---|--|--|
| Profile name                              | TEST9302   |  |
| Description                               | Z1DRB1 GA  |  |
| Load type                                 | <input type="radio"/> Normal <input checked="" type="radio"/> Clear <input type="radio"/> SCSI <input type="radio"/> SCSI dump |  |
| Load address                              | 9302   | <input type="checkbox"/> Use dynamically changed address   |
| Load parameter                            | C73001M1   | <input type="checkbox"/> Use dynamically changed parameter |
| Time-out value                            | 300  | 60 to 600 seconds  |
| <input type="checkbox"/> Store status     |  |  |
| Worldwide port name                       | 0  |  |
| Logical unit number                       | 0  |  |
| Boot program selector                     | 0  |  |
| Boot record logical block address         | 0  |  |
| Operating system specific load parameters |  |  |

Figure 7-51 HMC Load Profile

In Figure 7-51:

- The LOAD address points to where z/OS and programs reside
- The LOAD parameter C73001M1, where:
  - **C730** equals where the IODF resides and also the SYS#.IPLPARM dataset
  - **01** equals the suffix value of the LOADxx member in SYS#.IPLPARM
  - **M** equals the NIP suppression indicator
  - **1** equals alternate nucleus member

11. We must logon to the system where the IODF was written, in our example SC30, and access SYS#.IPLPARM.

12. Enter the command **D IPLINFO** to obtain which LOADxx member the system used during IPL. Refer to Figure 7-52.

```

-D IPLINFO
IEE254I 16.03.11 IPLINFO DISPLAY 432
SYSTEM IPLLED AT 22.03.35 ON 11/02/2011
RELEASE z/OS 01.13.00 LICENSE = z/OS
USED LOAD01 IN SYS0.IPLPARM ON OC730
ARCHLVL = 2 MTLSHARE = N
IEASYM LIST = XX
IEASYS LIST = (00) (0P)
IODF DEVICE: ORIGINAL(OC730) CURRENT(OC730)
IPL DEVICE: ORIGINAL(O9302) CURRENT(O9302) VOLUME(Z1DRB1)
  
```

Figure 7-52 MVS D IPLINFO command

Here is a sample of the LOAD01 member from SYS0.IPLPARM showing:

- HWNAME = SCZP301
- LPARNAME = A11

– IODF = 07

Example 7-7 shows HCD LOAD01 member in SYS0.IPLPARM.

*Example 7-7 HCD LOAD01 member in SYS0.IPLPARM*

---

```
*-----DEFINITION FOR SC30-----*
HWNAME  SCZP301
LPARNAME A11
*IEASYM 99
SYSPLEX COMPLEX Y
IODF    07 SYS6 TEST2094 01 Y
SYSCAT  COMCAT123CMCAT.COMCAT
PARMLIB SYS1.PARMLIB
PARMLIB CPAC.ZOSR1D.PARMLIB
PARMLIB SYS1.IBM.PARMLIB
*-----DEFINITION FOR SC31-----*
HWNAME  SCZP301
LPARNAME A13
*IEASYM 99
SYSPLEX COMPLEX Y
IODF    07 SYS6 TEST2094 01 Y
SYSCAT  COMCAT123CMCAT.COMCAT
PARMLIB SYS1.PARMLIB
PARMLIB CPAC.ZOSR1D.PARMLIB
PARMLIB SYS1.IBM.PARMLIB
```

---

This ensures that the operating system component of the IODF (software), when IPLed, matches the hardware component of the IODF (hardware) that is loaded into HSA.

## 7.6 Configuring the IBM Storage System DS8000

The host adapter ports in the DS8000 storage control unit must be configured for FICON (FC) protocol. Follow the procedure described in Appendix D, “Configuring the DS8000 for FICON” on page 415, to configure the host adapter ports.

For planning and implementation information, refer to *IBM System Storage DS8000: Architecture and Implementation*, SG24-8886.

### 7.6.1 Configuring the FICON Directors

The FICON Directors must be configured to match the designated configuration, as shown in Figure 7-1 on page 198. For the desired cascaded FICON Directors topology, this includes:

- ▶ Specifying the Domain ID in the Directors
- ▶ Enabling the ports
- ▶ Checking the port settings

Other settings in the FICON Director are optional and must be specified after reading the considerations explained in Chapter 4, “Planning the FICON environment” on page 77.

Make sure that the selected ports (LX/SX) in the Director match the FICON node (LX/SX) that is attached.

All of the required steps to configure the FICON Directors are described in detail in Chapter 8, “Configuring FICON Directors” on page 245. Go to 8.1.1, “Configuration flowchart” on page 246 and follow the procedures described there.

Return to this section after the FICON Directors are configured and ready for use.

## 7.6.2 Connecting the fiber optic cables

After all of the definition and configuration tasks are completed, fiber optic cables have to be plugged in to the FICON channels, CU host adapters, and the FICON Directors.

All the ports in our configuration are long wavelength (LX) laser types. For LX type lasers, a 9µm single mode (SM) fiber optic cable is recommended to achieve maximum performance and distance. All fiber optic cables must have an LC duplex connector at both ends.

Plug the fiber optic cable connectors to the designated ports (shown in Figure 7-1 on page 198) to the FICON channel, the CU, and the FICON Director, as listed here:

- ▶ PCHID 399 to FICON Director #65 port 21
- ▶ PCHID 3E1 to FICON Director #65 port A3
- ▶ ISL connection from FICON Director #65 port B4 to FICON Director #66 port 00
- ▶ ISL connection from FICON Director #65 port C6 to FICON Director #66 port 20
- ▶ CU host adapter port 0133 to FICON Director #66 port 41
- ▶ CU host adapter port 0333 to FICON Director #66 port 11

Make sure the fiber optic cable ends are cleaned before plugging them into the transceiver. Ensure that all fiber optic cables are labeled and documented for later use (for example, when recabling or for problem determination).

For further information and considerations regarding fiber optic cabling and documentation, refer to 4.11, “Physical connectivity” on page 113 and 4.2, “Documentation” on page 80.

## 7.6.3 Configuring the channel path and device online

After you successfully activate the new I/O definitions, you can configure the channel path online and try to communicate with the devices:

**Note:** Even if all the definitions in the IOCDs are correct, the channel status might be offline.

1. Enter **D M=CHP(21)** on the system console to display the current status of CHPID 21. If the CHPID is offline, enter **CF CHP(21),ON** to configure the channel online. The channel status changes to ONLINE, as shown in Example 7-8.

*Example 7-8 CHPID 21 online state*

---

```
- CF CHP(21),ONLINE
IEE502I CHP(21),ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

---

To achieve the best performance on the FICON channel, make sure that zHPF is enabled. Refer to 4.9.6, “High Performance FICON” on page 107 for considerations regarding how to exploit zHPF on System z servers.

2. Enter **D IOS,ZHPF** at the z/OS console to display zHPF settings for the z/OS image. If zHPF is disabled, enter **SETIOS ZHPF,YES** to enable zHPF, as shown in Example 7-9 on page 238.

Example 7-9 MVS - Displaying zHPF

-D IOS,ZHPF
IOS630I 13.27.23 ZHPF FACILITY 021
HIGH PERFORMANCE FICON FACILITY IS ENABLED

Using the SETIOS ZHPF=YES command enables zHPF temporarily. However, after the next system IPL, the zHPF facility is reset to the default (disabled). To permanently enable zHPF for z/OS, add or change the ZHPF parameter in the SYS1.PARMLIB member IECIOSxx to ZHPF=YES. Refer to Example 7-10.

Example 7-10 zHPF settings in SYS1.PARMLIB(IECIOS00)

EDIT SYS1.PARMLIB(IECIOS00) - 01.16 Columns 00001 00072
Command ===> Scroll ==> CSR
\*\*\*\*\* Top of Data \*\*\*\*\*
000001 HYPERPAV=YES
000002 ZHPF=YES
\*\*\*\*\* Bottom of Data \*\*\*\*\*

- 3. Now you can query the status and functional details of the channel by entering D M=CHP(21) at the operator console. The command output is shown in Example 7-11. It provides information about the channel and the attached devices.

Example 7-11 D M=CHP(21)

-D M=CHP(21)
IEE174I 16.17.36 DISPLAY M 521
CHPID 21: TYPE=1B, DESC=FICON SWITCHED, ONLINE
DEVICE STATUS FOR CHANNEL PATH 21
0 1 2 3 4 5 6 7 8 9 A B C D E F
0900 + + + + + + + + + + + + + + +
0901 + + + + + + + + + + + + + + +
0902 + + + + + + + + + + + + + + +
0903 \* \* \* \* \* \* \* \* \* \* \* \* \* \* \*
0904 \* \* \* \* \* \* \* \* \* \* \* \* \* \* \*
0905 \* \* \* \* \* \* \* \* \* \* \* \* \* \* \*
0906 \* \* \* \* \* \* \* \* \* \* \* \* \* \* \*
0907 + + \* \* \* \* \* \* \* \* \* \* \* \* \* +
....
1911 UL UL UL UL UL UL HA HA HA HA HA HA HA HA HA HA
1912 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1913 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1914 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1915 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1916 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
1917 HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
SWITCH DEVICE NUMBER = NONE
ATTACHED ND = SLKWRM.DCX.BRD.CA.1AFX0642C00Z
PHYSICAL CHANNEL ID = 0399
FACILITIES SUPPORTED = ZHPF

Example 7-11 shows that CHPID 21 is online and operating in a switched configuration. Information is also displayed about the attached devices and the facilities (for example: ZHPF) supported by the channel.

- 4. To verify that communication to the attached devices is working properly, enter D M=DEV(xxxx), where xxxx is any device number. For example, to check the status of

device 9000, which is a storage device, enter **D M=DEV(9000)** on a z/OS console (see Example 7-12).

*Example 7-12 MVS D M=DEV(9000) with device offline*

---

```
-D M=DEV(9000)
IEE174I 16.23.36 DISPLAY M 539
DEVICE 9000 STATUS=OFFLINE
CHP                40  41  42  43  21  25  22  26
ENTRY LINK ADDRESS  22  22  1E  1E  6592 65A3 6500 6519
DEST LINK ADDRESS  1B  1B  10  10  6641 6611 6534 6562
PATH ONLINE        Y   Y   Y   Y   Y   Y   Y   Y
CHP PHYSICALLY ONLINE Y   Y   Y   Y   Y   Y   Y   Y
PATH OPERATIONAL   Y   Y   Y   Y   Y   Y   Y   Y
MANAGED            N   N   N   N   N   N   N   N
CU NUMBER          9000 9000 9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED      = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED     = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW
```

---

**Note:** Observe that the ZHPF function is not displayed for this device. We found in our tests that after the device was brought online zHPF showed as enabled. Refer to Example 7-13.

*Example 7-13 MVS D M=DEV(9000) with device online*

---

```
-R0 SC30,D M=DEV(9000)
IEE174I 16.50.44 DISPLAY M 686
DEVICE 9000 STATUS=ONLINE
CHP                40  41  42  43  21  25  22  26
ENTRY LINK ADDRESS  22  22  1E  1E  6592 65A3 6500 6519
DEST LINK ADDRESS  1B  1B  10  10  6641 6611 6534 6562
PATH ONLINE        Y   Y   Y   Y   Y   Y   Y   Y
CHP PHYSICALLY ONLINE Y   Y   Y   Y   Y   Y   Y   Y
PATH OPERATIONAL   Y   Y   Y   Y   Y   Y   Y   Y
MANAGED            N   N   N   N   N   N   N   N
CU NUMBER          9000 9000 9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED:  0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED      = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED     = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW, ZHPF
```

---

The response to **D M=DEV(9000)** displays all available paths to the devices and their status. Information about the device (for example, device type), the control unit number, and the functions supported by the device (for example, MIDAW and ZHPF) is shown.

## 7.6.4 Verifying the installation

After all of the configuration tasks in the z196 server and the DS8000 storage are complete, you can verify that the current configuration matches the desired configuration shown in Figure 7-1 on page 198.

Perform this verification to ensure that the fiber optic cables are correctly plugged to the designated ports in the System z196 server and the DS8000, and to check the status of channels and devices.

To verify installation:

1. Check the status of PCHID 399 in the z196 server. From any HMC where the desired CPC (server) is defined (for example, SCZP301), start a Single Object Operation (SOO) to that CPC's Support Element.

Display the list of channels installed on the CPC. Search the list for PCHID (399), and double-click the PCHID icon. The PCHID details panel shown in Figure 7-53 is displayed.

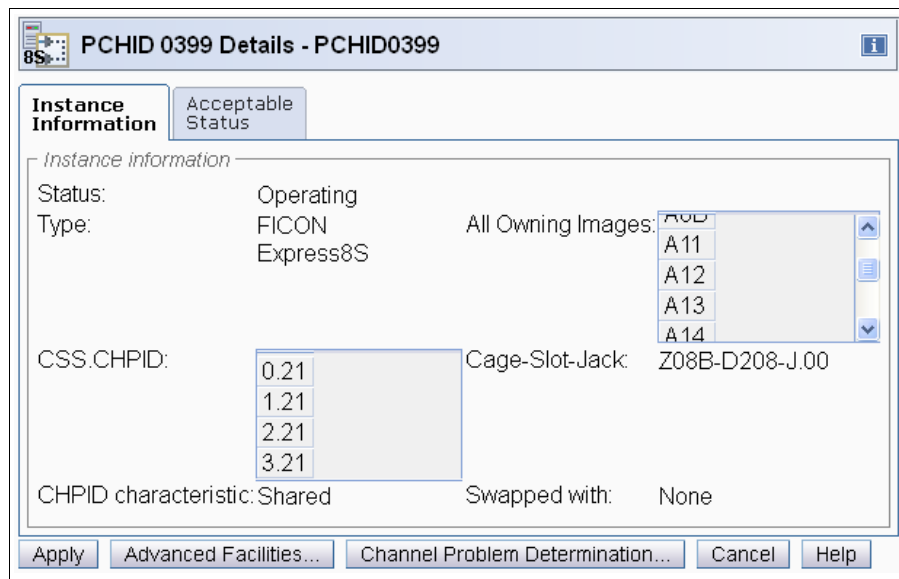


Figure 7-53 HMC PCHID details for 399

Essential information for PCHID 399 is shown on the PCHID Details panel:

- The PCHID status is Operating
- The PCHID type is FICON Express8S
- CSS.CHPID 1.21 is assigned to PCHID 399 (plus others)
- The owning images of PCHID 399 are A11 (plus others)
- The CHPID assigned to PCHID 399 is shared

2. The CHPID Details panel provides information that is similar to the PCHID details panel. To display the CHPID details, select the channel list for an LPAR from the Support Element. Refer to Figure 7-54 on page 241.






Figure 7-54 HMC CHPID details for 1.21

Notice that the information for PCHID 399 (CHPID 1.21) provided on the details window matches the designated configuration. This proves that PCHID 399 has CHPID 1.21 assigned and that image A11 can access the channel. Repeat these checks on other channels that were recently configured.

**Important:** If any of the data displayed on the PCHID or CHPID detail panel does not match the desired configuration, you must correct the definitions in the IOCDS.

3. Check that the FICON channels are connected to the correct FICON switch and port. On either the PCHID details or CHPID details panel, click **Channel Problem Determination** to display the Channel Problem Determination panel, where you can select which information you want to display (see Figure 7-55).

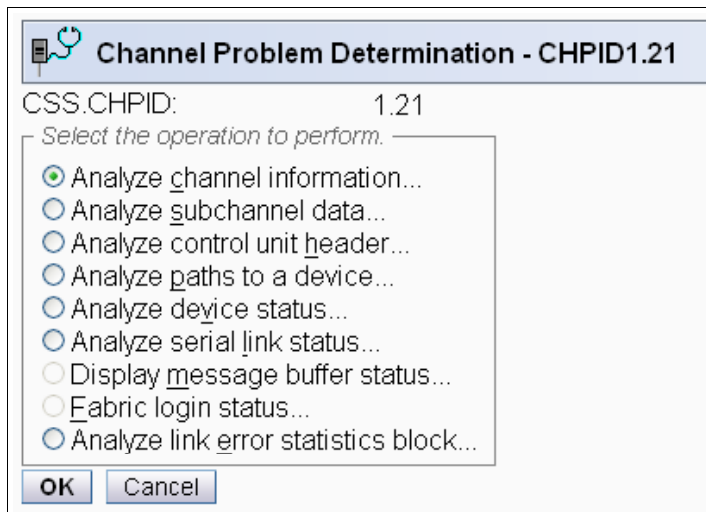


Figure 7-55 HMC Channel Problem Determination panel

4. Select **Analyze channel information**, and click **OK**. The Analyze Channel Information window is displayed, as shown in Figure 7-56, which provides information about the node attached to the FICON channel.

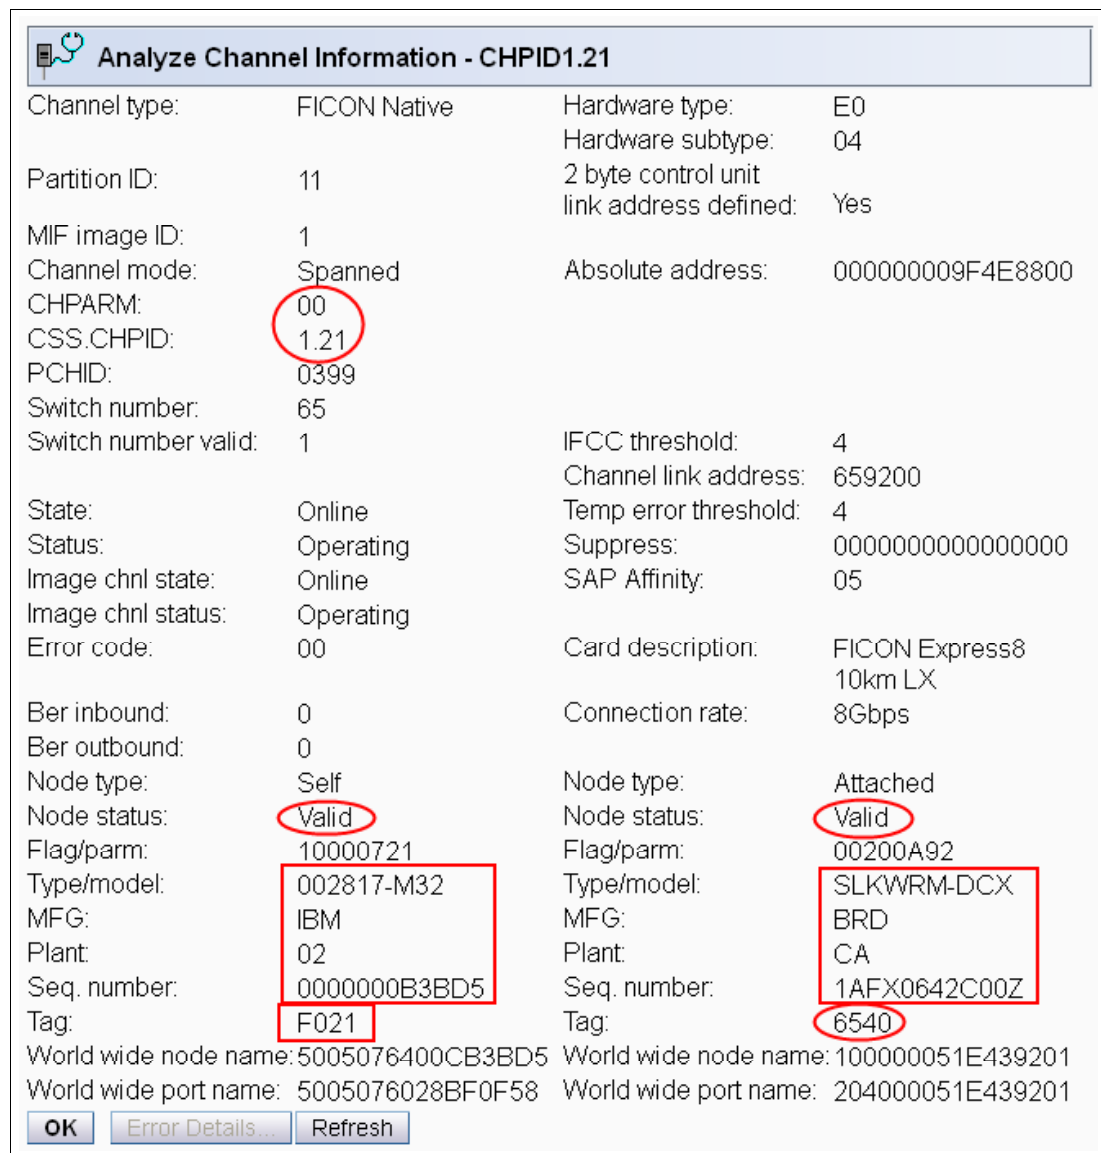


Figure 7-56 HMC Analyze Channel Information

Information about the nodes is displayed at the bottom part of the window. The lower-left side displays information about the node in the z196 server. The lower-right side displays information about the attached node.

**Important:** Make sure that the Node status for both nodes is displayed as **Valid**. If any other status is shown, none of the displayed information is valid.

5. Check that the Type/model information and the serial number (Seq. number) is as expected.
6. Examine the Tag field for each node. The Tag provides information about the port number of the attached node. For the z196 server, the two right-most digits of the Tag value represent the CHPID number for the channel node (21). For the attached node, the four

digits represent the switch port number (6540). Be aware, however, that the tag value is provided by the attached device during link initialization and might have different meanings, depending on the vendor.

- The World Wide Node Name (WWNN) and the World Wide Port Name (WWPN) are also shown for each port, and can be used to prove that the channel is connected to the correct FICON adapter port if the WWNN or WWPN of the attached device is known.

If the node status is not Valid or the Tag value and WWPN value are not correct, check the fiber optic cable link between the z196 server and the FICON Director or DS8000 host adapter to ensure that it is plugged correctly.

In our scenario, we are now sure that PCHID 399 has connectivity to switch 65 and port 92, which matches our desired configuration (see Figure 7-1 on page 198).

If the displayed values are not as expected, the fiber optic cables might not be plugged correctly and must be checked.

After completing the preceding steps and proving that the physical path to the DS8000 storage controller and the logical definitions of the link are correct, check that the path to the control unit image is initialized correctly and properly defined.

- On the Channel Problem Determination panel, shown in Figure 7-55 on page 241, select **Analyze Serial Link Status**, and click **OK**. The Analyze Serial Link Status window is displayed, as shown in Figure 7-57.

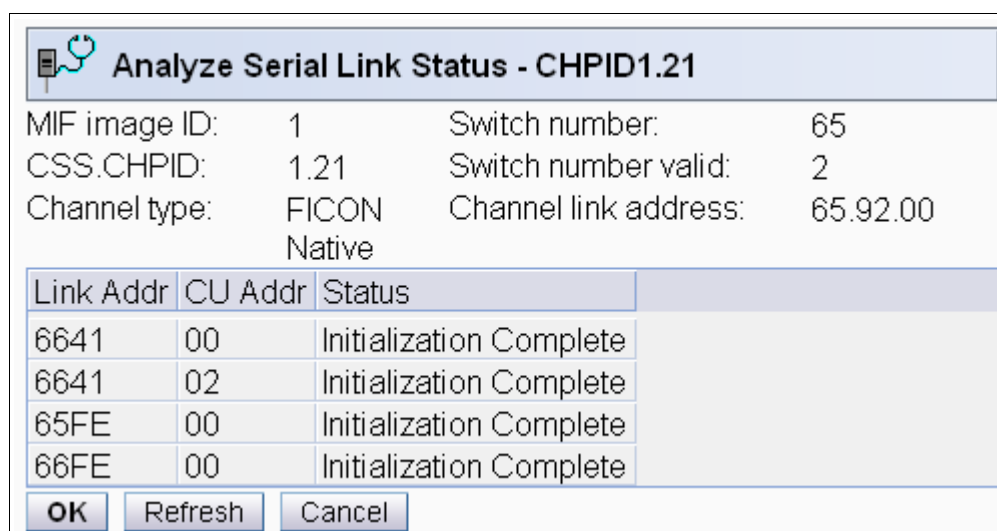


Figure 7-57 HMC Analyze Serial Link status

The Analyze Serial Link Status window provides status information about the link to the control unit images defined in the IOCDs. Scroll through the list of CU images and check that status for all CUs is Initialization Complete. In this case, Figure 7-57 shows a link status of **Initialization Complete** for all defined CU images on link address 65FE, 66FE and 6641.

In a cascaded FICON configuration, you must use a two-byte link address to specify the link address. That means link address 65FE is the CUP port in switch 65, and link address 6641 points to a DS8000 storage control unit attached to switch 66.

Link address 6641 is the destination port where the DS8000 control unit port is physically attached. The link to CU address 00 and 02 is initialized. CU address 00 and 02 correspond to control unit number 9000 (CUADD=0) and 9100 (CUADD=2), as defined in the IODF.

If the link status **Initialization Complete** is not shown, you must check that the ports in the FICON Director and CU are correctly configured, and that the fiber optic cable link between the Director and the CU has the correct cable type and plugging.



## Configuring FICON Directors

In this chapter, we explain how to configure IBM 2499 FICON Directors. We provide configuration examples for switched and cascaded FICON Director topologies.

The following topics are discussed:

- ▶ Configuration overview
- ▶ Installing and using IBM Network Advisor
- ▶ Setting up a FICON Director
- ▶ Setting up a cascaded topology
- ▶ FICON Directors in an extended distance environment
- ▶ FICON Directors in an intermixed environment
- ▶ Channel swap
- ▶ Backing up Director configuration data
- ▶ Backing up IBM Network Advisor configuration data

For information about configuring Cisco MDS FICON Directors, refer to *Cisco FICON Basic Implementation*, REDP-4392.

## 8.1 Configuration overview

We show our approach to configuring FICON Directors in different environments. All examples are based on an IBM System Storage SAN768B-2 Director (2499-816) and an IBM System Storage SAN384B-2 Director (2499-416). Configuration for all other B-type products is similar. The window captures are based on IBM Network Advisor Version 11.1.2 and on Fabric OS (FOS) Version 7.0.0.c

For switched and cascaded FICON Director environments, we illustrate how to set up:

- ▶ IP addresses (for switch management)
- ▶ Licensed features
- ▶ Logical switches
- ▶ Domain ID and Insistent Domain ID
- ▶ Port-Based Routing (PBR), Dynamic Load Sharing (DLS) and In-Order Delivery (IOD)
- ▶ Control Unit Port (CUP)
- ▶ Port type and speed
- ▶ Buffer credits
- ▶ Allow/Prohibit matrix
- ▶ Zoning
- ▶ Port fencing
- ▶ Inter-Switch Links (ISLs) including Traffic Isolation (TI) and Quality of Service (QoS)
- ▶ High-integrity fabric and policies (fabric binding)

The configuration described in this chapter is based on considerations explained in Chapter 4, “Planning the FICON environment” on page 77.

**Terminology:** Throughout this chapter, the terms FICON Director and switch are used interchangeably. In all cases, the term switch refers to FICON Director.

### 8.1.1 Configuration flowchart

The flowchart displayed in Figure 8-1 on page 247 illustrates the approach we used to configure our FICON Directors for the environments discussed in Chapter 6, “Configuring a switched topology” on page 155 and Chapter 7, “Configuring a cascaded topology” on page 197.

Configuring switched and cascaded topologies are similar. The only differences are the link addressing scheme on the System z server side, the need for the high-integrity fabric function, and the implementation of ISLs; therefore, 8.3, “Setting up a FICON Director” on page 262 applies to both topologies. Section 8.4, “Setting up a cascaded topology” on page 298 only applies to cascaded FICON Directors.

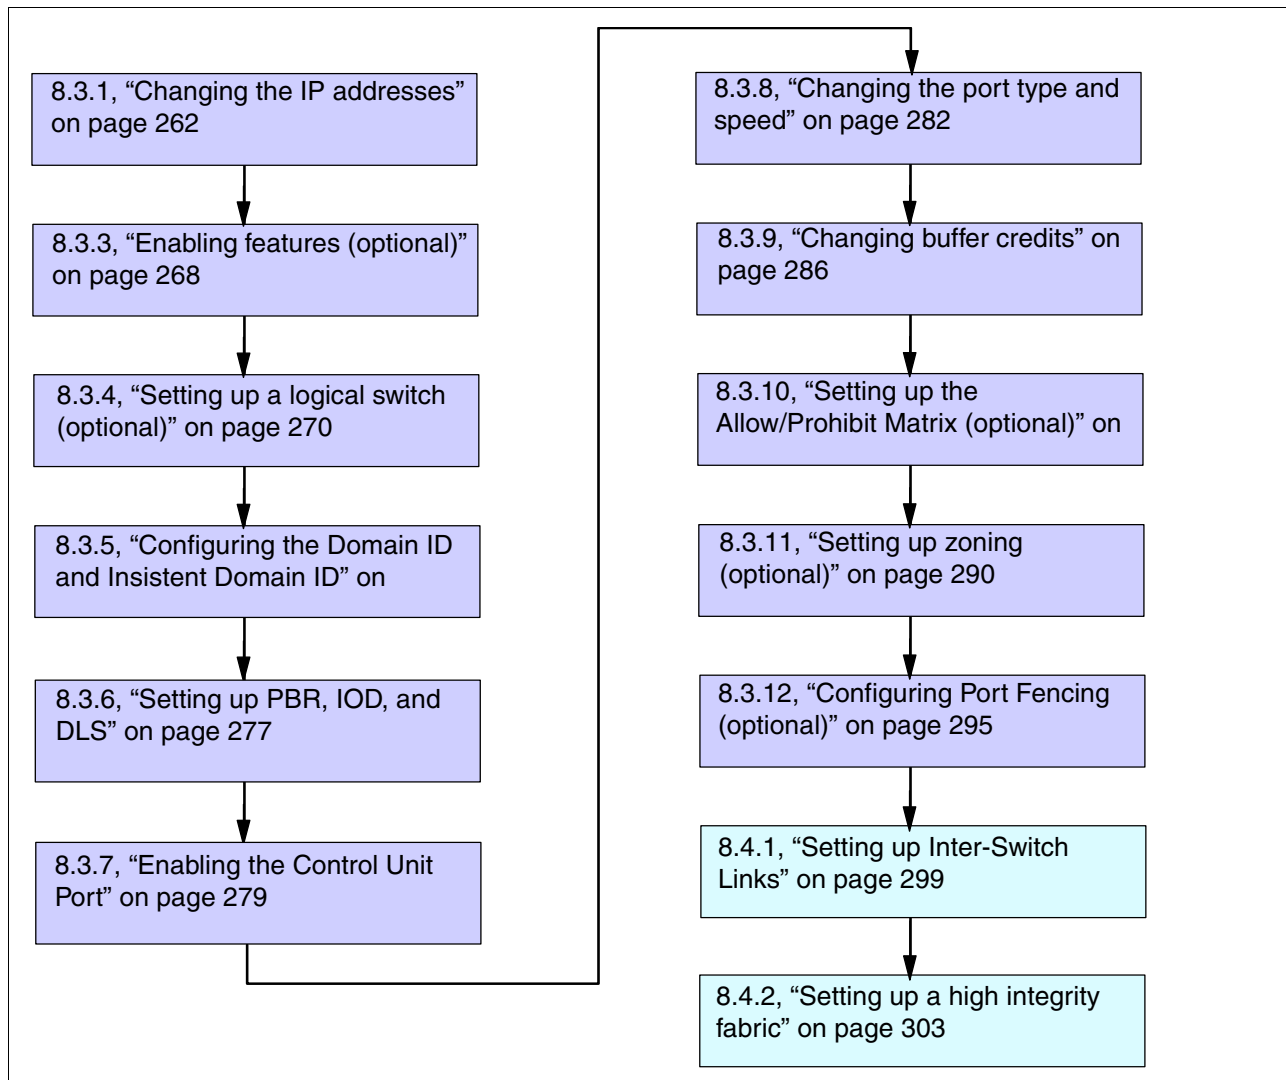


Figure 8-1 Configuration flowchart

Depending on your requirements, some optional steps might not be needed, and the last two steps are only relevant to a cascaded topology. However, if you plan to install a single FICON Director and want to use the 2-byte link addressing on the System z server for future use in a two-site configuration, you must configure the Switch Connection Control (SCC) as described in 8.4.2, “Setting up a high integrity fabric” on page 303. In the SCC policy, you only have one WWNN entry (from that particular FICON Director).

Before we started with the configuration of the FICON Directors, we installed the IBM Network Advisor. IBM Network Advisor was used to set up our FICON Directors and to configure some optional functions.

## 8.1.2 FICON Director management connectivity

The FICON Directors and IBM Network Advisor servers must always be interconnected on an isolated LAN to isolate Director management traffic from other IP traffic.

For controlling and maintaining the IBM 2499 FICON Directors, both CP cards in each chassis must have connectivity to the IBM Network Advisor server. For each FICON Director,

a total of three IP addresses are required: two for the CP cards, and one for the Chassis. One IP address is required for each IBM Network Advisor server.

**Note:** IP addresses in the range of 10.0.0.1 to 10.0.0.255 are reserved for the internal communication of the IBM 2499 FICON Directors and must not be used for the CP cards or Directors.

Because the IBM Network Advisor server does direct polling of the fabric to gather configuration and status information, they must be able to access every Director in the fabric. Any Director that IBM Network Advisor server cannot reach goes grey within their respective IBM Network Advisor application.

Our FICON Director management environment is shown in Figure 8-2. It also shows the IP addresses used for our FICON Directors, IBM Network Advisor server, and firewall/router.

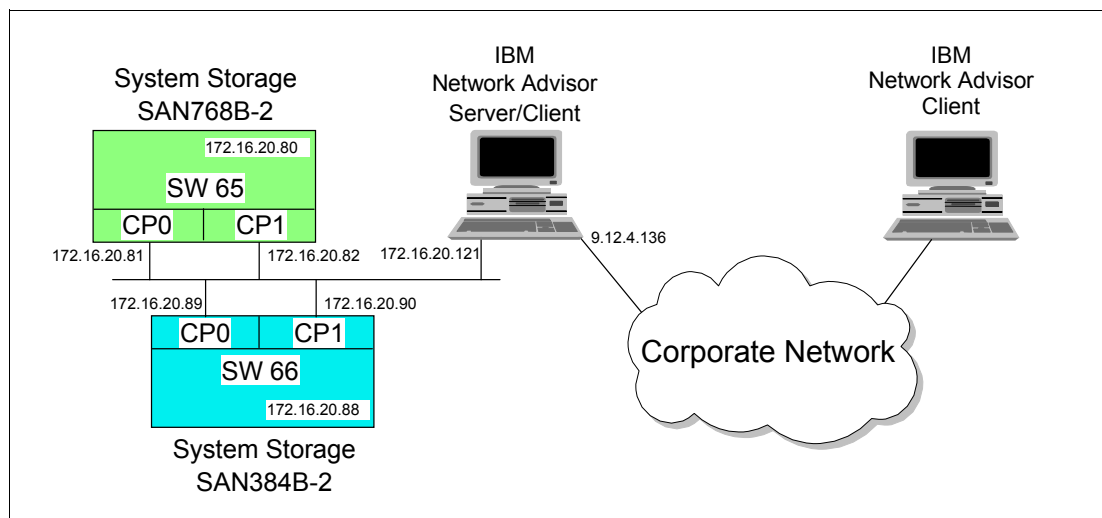


Figure 8-2 Our FICON Director management environment

In this configuration, the IBM Network Advisor Server participates in two Ethernet networks: the Service LAN connected to all CP cards in the FICON Directors and the Corporate LAN.

## 8.2 Installing and using IBM Network Advisor

The IBM Network Advisor software is used to manage several Directors and switches through one graphical user interface (GUI). With the IBM Network Advisor you can configure and monitor all fabrics, Directors, and switches that are installed and connected through the network to the Management Server.

The IBM Network Advisor is a client-server management software package. After the server software is installed, all other users can download the client part of the package by typing the server's IP address in a web browser address field.

Check the IBM Network Advisor data sheet for hardware and software prerequisites at:

[http://www-01.ibm.com/common/ssi/rep\\_ca/3/897/ENUS211-243/ENUS211-243.PDF](http://www-01.ibm.com/common/ssi/rep_ca/3/897/ENUS211-243/ENUS211-243.PDF)



Review the release notes at:

[ftp://public.dhe.ibm.com/storage/san/networkadvisor/IBM\\_NetworkAdvisor\\_v11.1.1\\_Release\\_Notes.pdf](ftp://public.dhe.ibm.com/storage/san/networkadvisor/IBM_NetworkAdvisor_v11.1.1_Release_Notes.pdf)

## 8.2.1 Installing IBM Network Advisor

To install IBM Network Advisor on the server:

1. Get the installation DVD or the installation package for IBM Network Advisor. If autorun is enabled, the installation will start, as shown in Figure 8-3; otherwise, double-click the installation package from the installation source (located in folder Windows\install.exe).

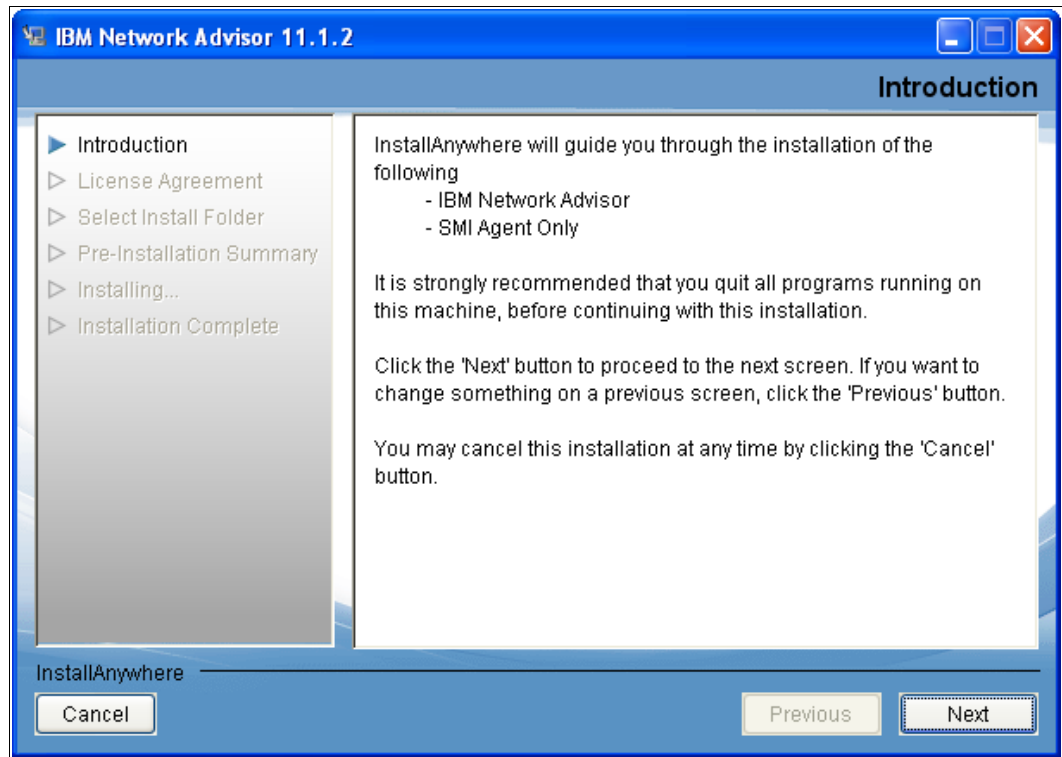


Figure 8-3 IBM Network Advisor Introduction panel

2. Follow the instructions shown in Figure 8-3 until you reach the Installation Complete window shown in Figure 8-4 on page 250.

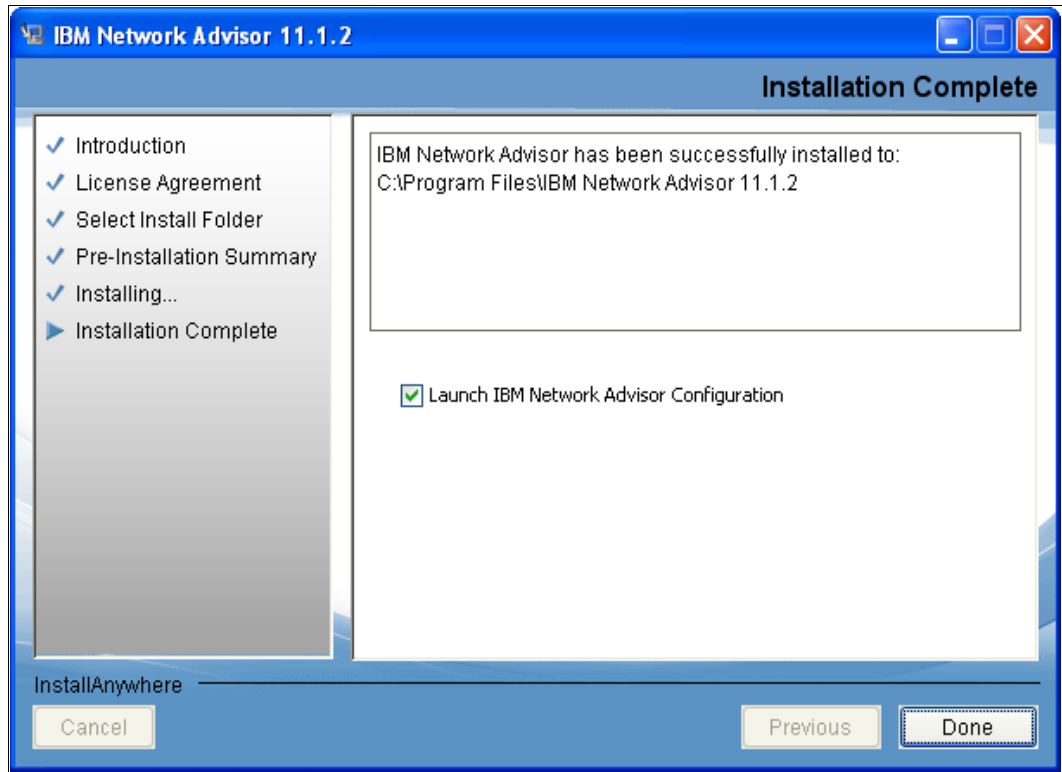


Figure 8-4 Installation Complete window

Ensure that **Launch IBM Network Advisor Configuration** is selected. Click **Done**.

3. On the Welcome window, click **Next**.
4. If you do *not* need to migrate data from a previous installation, select **No**, and click **Next** (see Figure 8-5 on page 251).

If you are migrating data, select the appropriate field, and then type the location of the installation you want to import. You can import data from a Data Center Fabric Manager, Enterprise Fabric Connectivity Manager (EFCM), a Fabric Manager (FM), or from a previous IBM Network Advisor installation.

**Note:** You must migrate any data to IBM Network Advisor during installation. You cannot migrate after installation time.

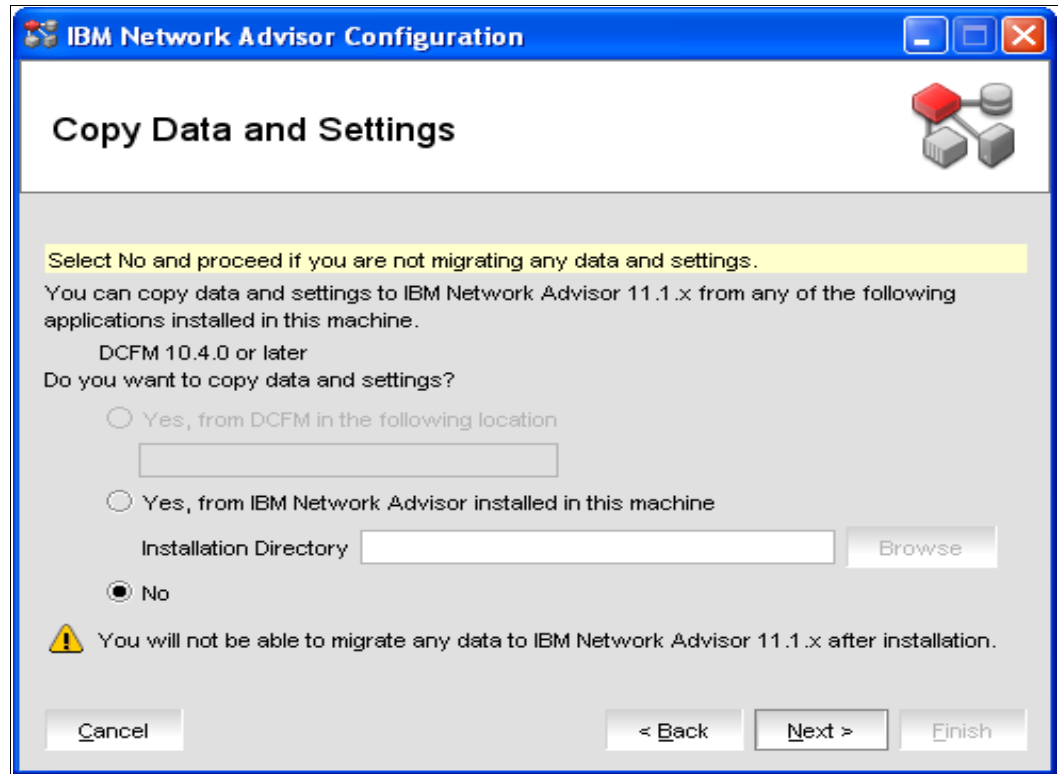


Figure 8-5 Data migration window

5. On the Package window, select SAN with SMI agent, as shown in Figure 8-6, and click **Next**.

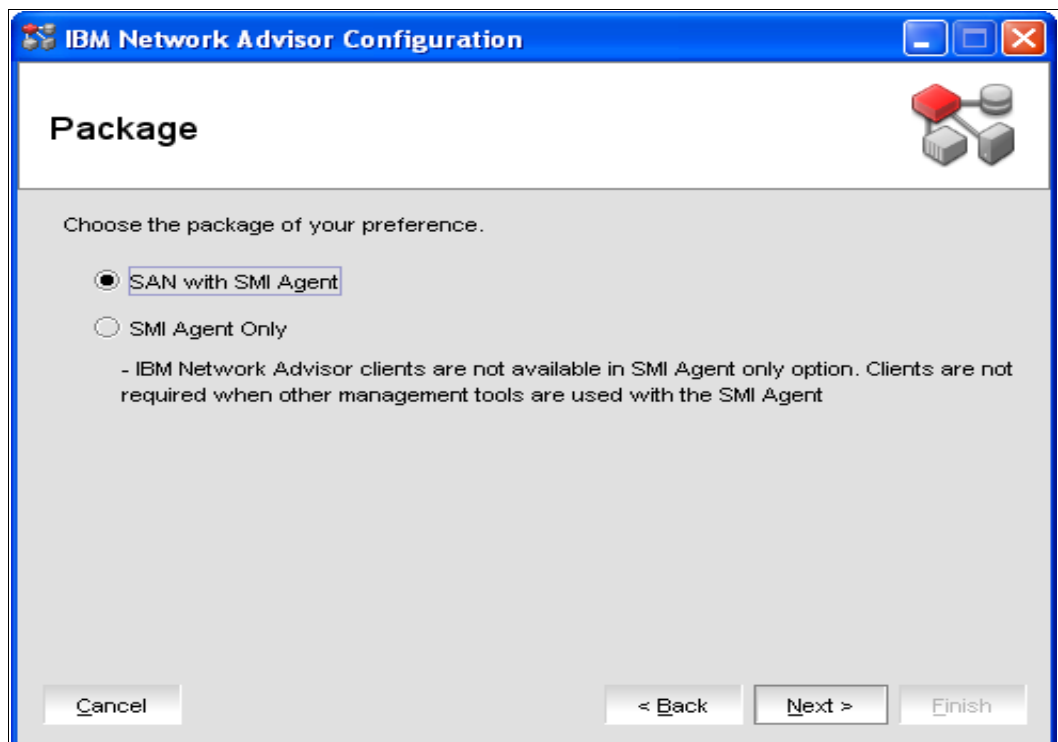


Figure 8-6 Select Package

**Note:** IBM Network Advisor clients are not available in SMI Agent only option. Clients are not required when other management tools are used with SMI Agent.

6. Select IBM Network Advisor - Licensed version, as shown in Figure 8-7, and click **Next**.

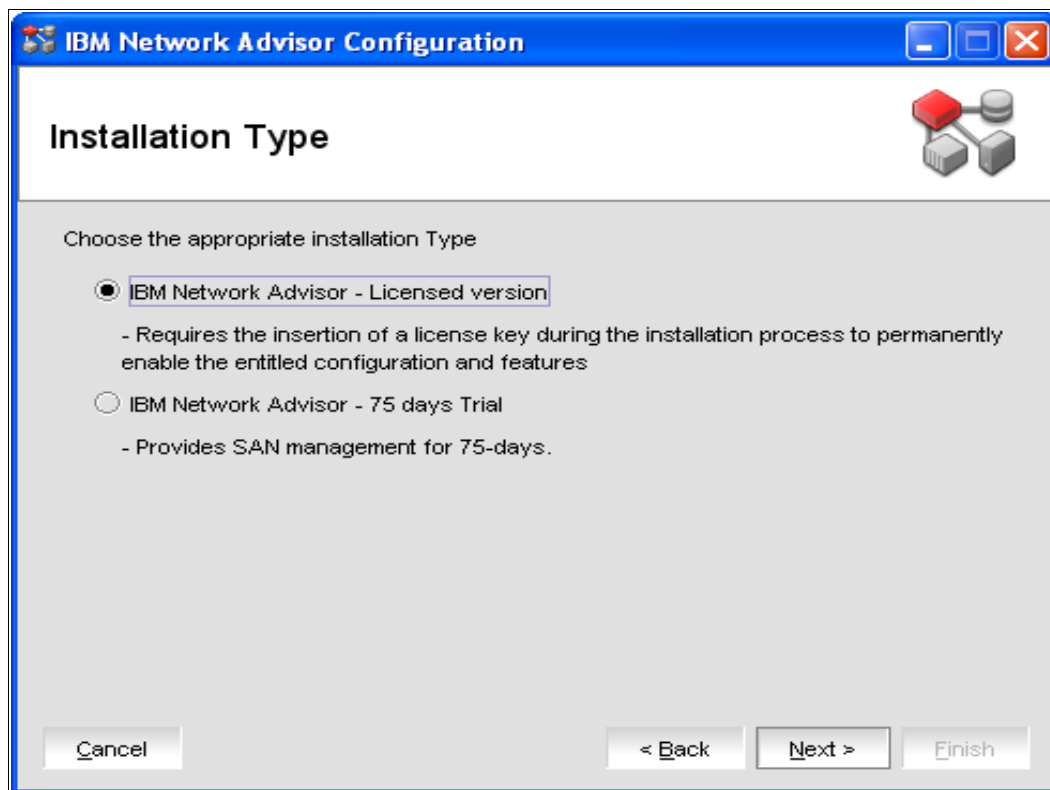


Figure 8-7 Installation type

7. Enter the serial number (from the DVD box) and the Server License (from the Key Certificate), or browse for and select the license file. If the window does not appear, your installation does not require a license. The License Key field is not case sensitive. See Figure 8-8 on page 253, and then click **Next**.

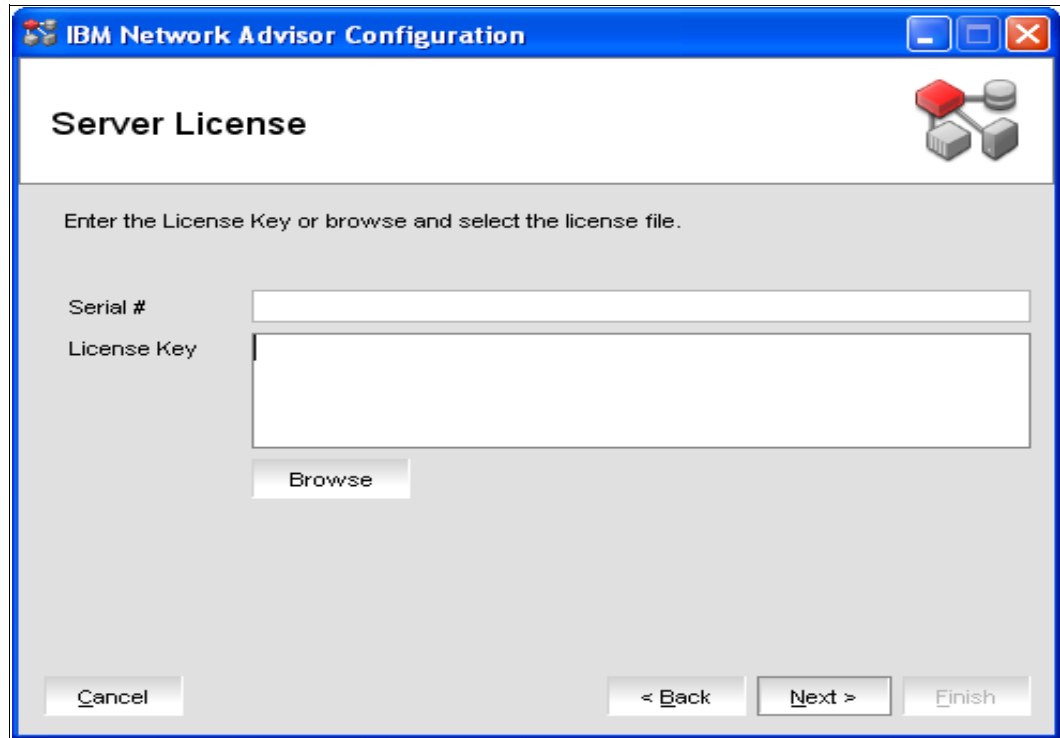


Figure 8-8 Server License key

8. Select Internal FTP Server, as shown in Figure 8-9, and click **Next**.

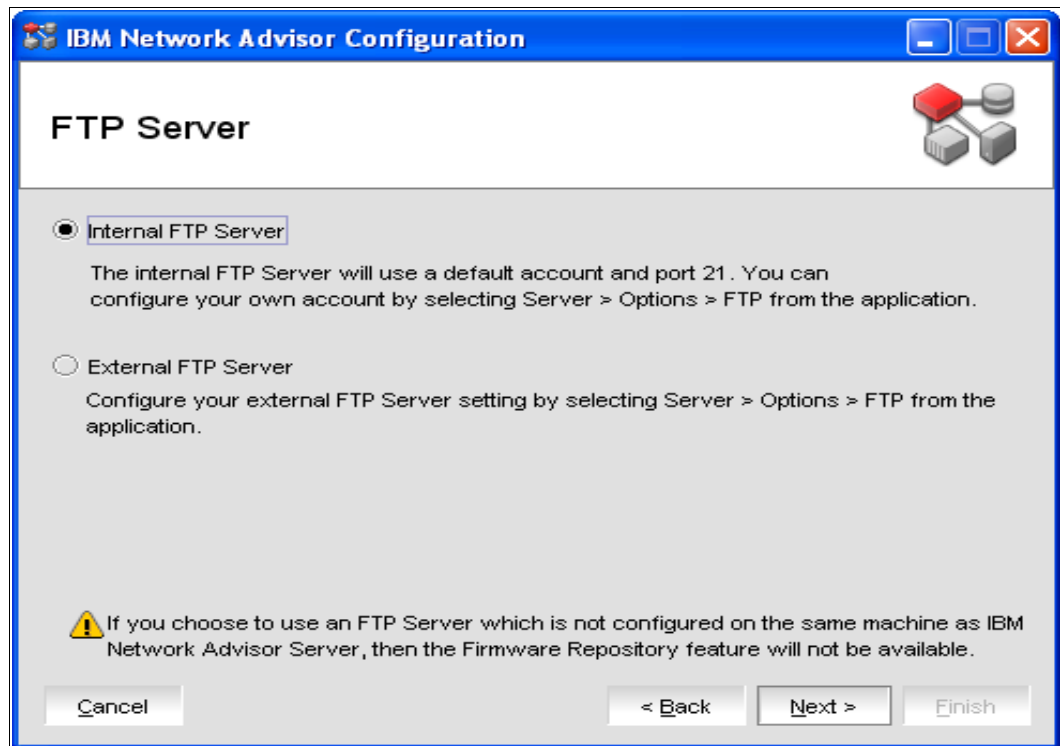


Figure 8-9 Select FTP Server

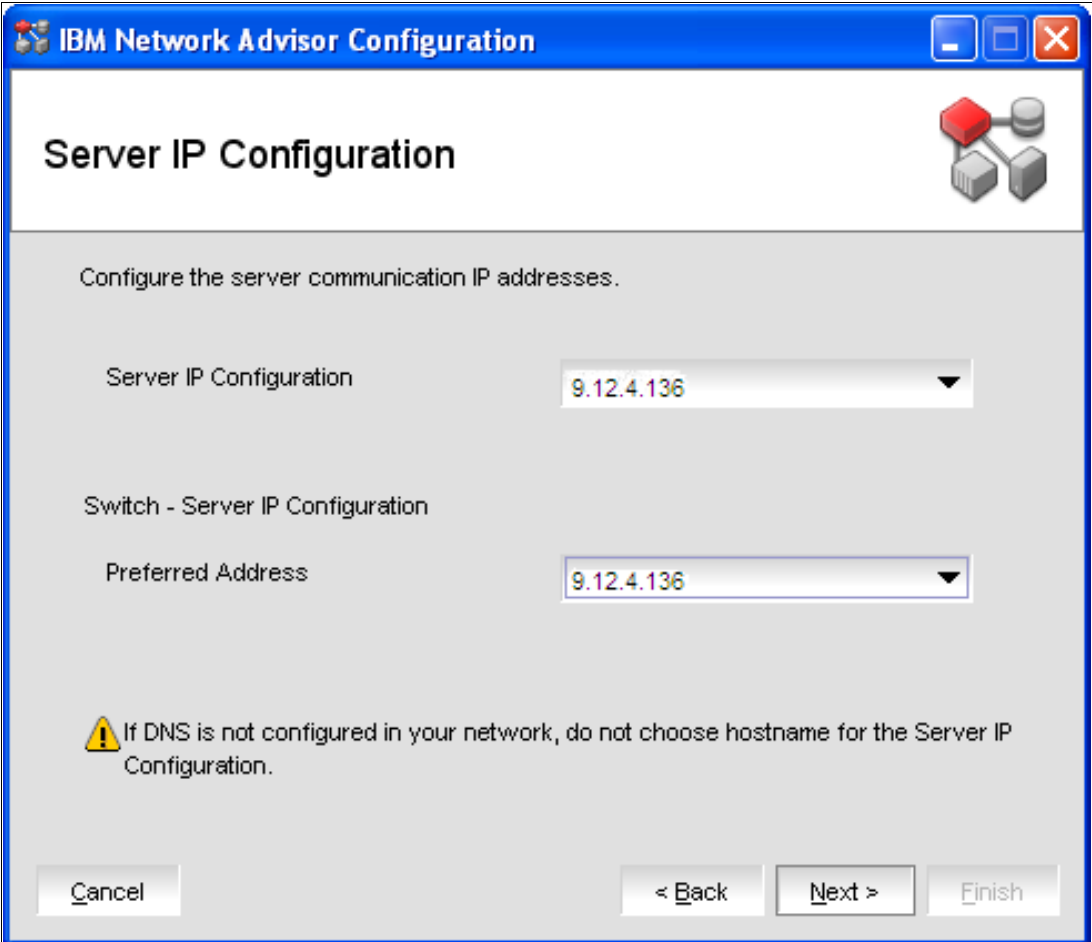
**Note:** You can change the FTP Server settings later in the IBM Network Advisor by selecting **SAN** → **Options**. In the Options window, select **Software Configuration** → **FTP/SCP** and make the required changes. Unless the ramifications of using an external FTP server are fully understood, always select the Internal FTP Server option.

9. From the pull-down menu in Figure 8-10, select the IP address that will be used by the clients as Return Address and the address that will be used to connect to the Directors and switches as Preferred Address. If only one network adapter is installed in the server, then the address will be the same for both pull-down menus.

**Important:** Do *not* select 127.0.0.1 as Return Address because if it is selected, clients cannot connect to the server's web application to download the client part of the software.

If you prefer to select the server name, make sure that it is configured in the DNS server.

10. After making your selection, as shown in Figure 8-10, click **Next**.



The screenshot shows a window titled "IBM Network Advisor Configuration" with a sub-header "Server IP Configuration". The main content area contains the instruction "Configure the server communication IP addresses." Below this are two dropdown menus: "Server IP Configuration" and "Switch - Server IP Configuration", both displaying the IP address "9.12.4.136". At the bottom, there is a warning icon and the text: "If DNS is not configured in your network, do not choose hostname for the Server IP Configuration." Navigation buttons at the bottom include "Cancel", "< Back", "Next >", and "Finish".

Figure 8-10 IBM Network Advisor Server IP address configuration

11. Enter the appropriate port numbers (see Figure 8-11), if you are not using the default values. The Syslog Port # must remain as 514 because this port number cannot be changed in the FICON Director.

**Note:** If a firewall is installed between the IBM Network Advisor server or client and the FICON Directors, the configured ports must be defined in the firewall to allow that traffic to flow.

Keep in mind that 16 consecutive ports are needed from the defined Starting Port #. Also, do not use port 2638 because it is used internally by the IBM Network Advisor server.

IBM Network Advisor Configuration

### Server Configuration

IBM Network Advisor requires Web Server, Database, Syslog and SNMP port numbers, as well as 16 consecutive port numbers from a Starting port #.

SSL Enabled

Web Server Port #

Database Port #

Starting Port #

Syslog Port #

SNMP Port #

Change this configuration by selecting Server > Options > Server Port from the application.

Figure 8-11 Configure network ports

**Tip:** You can change these port numbers later in the IBM Network Advisor by selecting **SAN** → **Options**. In the Options window, select **Software Configuration** → **Server Port** to make the changes.

To change the SNMP Port #, select **Monitor** → **SNMP Traps**.

12. Enable SMI Agent, as shown in Figure 8-12 on page 256. (Enable SLP and Enable SSL are only available after you select the **Enable SMI Agent** option), and click **Next**.

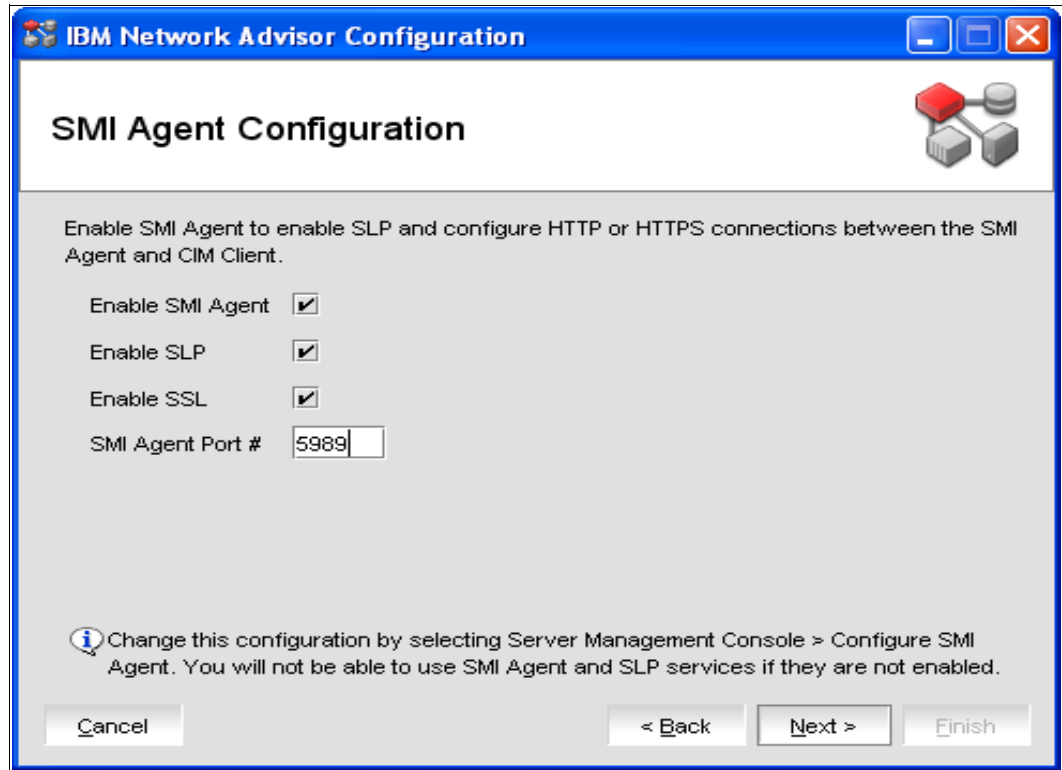


Figure 8-12 SMI Agent Configuration

13. Click **Next**, and select the size of your FICON/SAN installation, which is managed by the IBM Network Advisor server, as shown in Figure 8-13 on page 257.



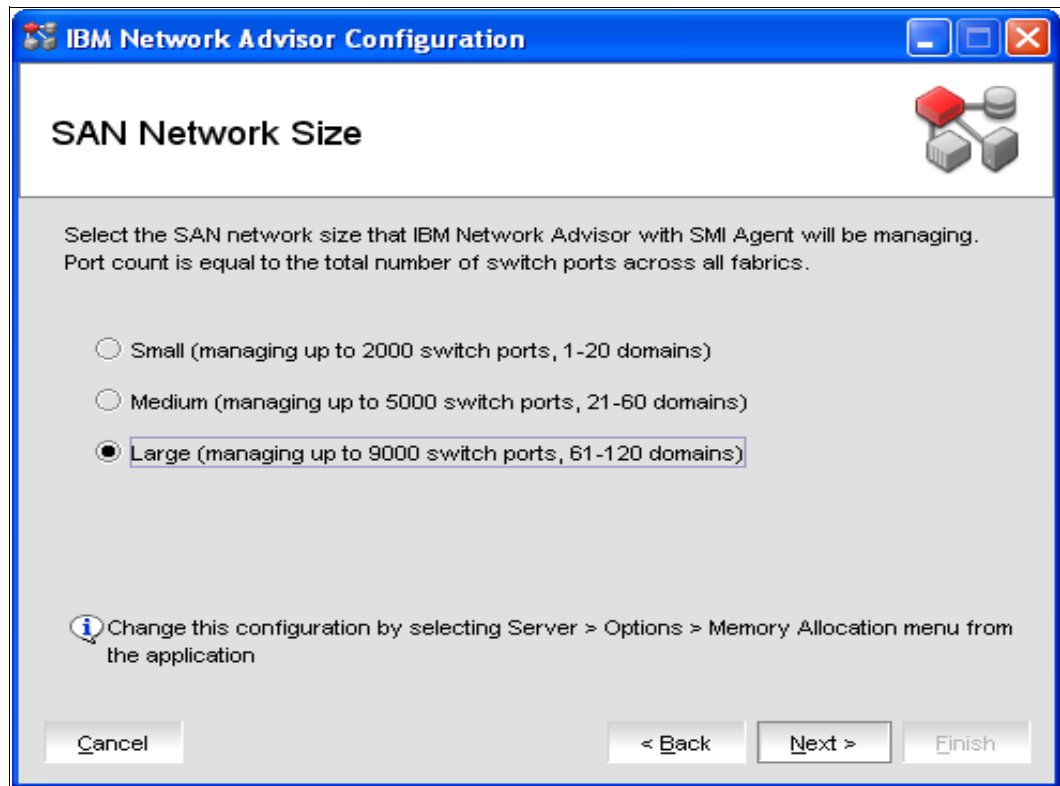


Figure 8-13 Select SAN size

**Note:** To ensure the fastest updates in Network Advisor, the recommended best practice is to select the smallest SAN network size that meets your requirements.

14. Click **Next** after you make the selection.
15. Wait until the database is initialized, and then click **Next** at the Server Configuration Summary window.
16. At the Start Server window, select **Start IBM Network Advisor Client**, and then click **Finish**.
17. Wait until the login window appears, shown in Figure 8-14 on page 258. If this is a migration or upgrade from a previous management software installation, you must use the user and password from that version. If it is a new install, the user is administrator and the password is password. Click **Login**.



Figure 8-14 IBM Network Advisor Login window

18. You will see the IBM Network Advisor Login Banner, which you can change or remove if it is not needed. This can be done after clicking **OK**. Figure 8-15 will appear.




Figure 8-15 IBM Network Advisor first startup

19. Select **SAN** → **Options**. In the Options window, select **Security Misc.**, and change the text in the Banner Message field or just remove the check mark, if it is not needed.

## 8.2.2 Using the IBM Network Advisor

In this section, we explain functions in the IBM Network Advisor that are needed during setup. For other functions, you can always use the Help button in the IBM Network Advisor.

### Discovering a FICON Director

Before you can manage a FICON Director, it must be discovered:

1. Select **Discover** → **Fabric** in the menu bar of the IBM Network Advisor.
2. Click **Add** to add a Director. A new window appears.
3. Enter the Director's chassis IP address, username (admin), and the password (the default is password). You can also provide a Fabric Name.
4. Select the **SNMP** tab, and select **v3** in the SNMP Version drop-down menu. Accept the default values, and then click **OK**.
5. Click **Close** to leave the discovery dialog.

The FICON Director is now discovered and can be managed through IBM Network Advisor.

If you want to change the admin password, select **Edit**, and change the password for the selected Director. Click **OK** after the password is changed. (Note that this password change is only for the communication between IBM Network Advisor and the Director. The password for the user ID must match the password set for the user ID on the chassis. To change the passwords on the chassis refer to "Creating and changing users and passwords".)

### Rediscovering a FICON Director

During the configuration, it might happen that you must rediscover a Director. To do this:

1. Select the discovered Director, and click **Delete**. The Director shows up in the lower part of the Discover Setup window.
2. Select the Director, and click **Discover** to rediscover it.
3. Enter the new Fabric Name, and click **OK**.

### Using the IBM Network Advisor client

To use the client from any personal computer, you must open a web browser, and enter the IP address of the IBM Network Advisor server in the web browser's address field. You will reach the welcome window from the IBM Network Advisor server.

You can also download the required Java Version and the MIB tables, if needed, for external system monitoring software. Click **Web Start the IBM Network Advisor Client**. The client is downloaded to your workstation and then the login window will appear. The handling is exactly the same as at the server. A shortcut is created in your program list so that you do not need to download the client again.

### Creating and changing users and passwords

To change the password, or create a new user for the IBM Network Advisor:

1. Select **SAN** → **Users** at the IBM Network Advisor menu bar.
2. Click **Edit** at the Users field if you want to change the password.
3. Click **OK** after the password is changed.

To create a new user, click **Add** and assign the user to a Resource Group by clicking the arrow pointing to the right. You can also create new user roles or edit an existing user role at

the Roles field. If you click **Add** in the Resource Groups field, you can build a new group that includes only the Directors that you assign to it. Click **OK** to activate all your changes.

To change the passwords on the Director:

1. Right-click the Director in the IBM Network Advisor, and select **Element Manager** → **Admin**.
2. In the Switch Administrator window, click the **User** tab.
3. Click **Change Password** to change the password for the selected user.
4. Click **OK** after the password is changed.
5. Click **Apply** to activate the changes.

To change the Open Database Connectivity (ODBC) password on the IBM Network Advisor server, go to the IBM Network Advisor program directory in the command prompt, and enter the following command `dbpassword <Username> <Password> <New_Password> <Confirm_Password>`, as shown in Example 8-1.

*Example 8-1 Changing the ODBC password*

---

```
C:\Program Files\IBM Network Advisor 10.1.3\bin>dbpassword guest password
test1test test1test
DB is updated Successfully
C:\Program Files\IBM Network Advisor 10.1.3\bin>
```

---

Now you must uninstall the ODBC driver on all workstations that use this function and install it again with the new password.

## Renaming the Director and ports

To change the Director name, right-click the Director in the IBM Network Advisor and select **Properties**. The properties window opens. Double-click in the Name field and type the new name. You can also change the port names by clicking the Port tab in the properties window and clicking in the Name field to change the port name. Click **OK** after all changes are done.

The port name can also be changed by right-clicking the Director and selecting **Element Manager** → **Ports**. Select the Port in the Port Administration window and click the **Rename** button. Type the name for the port and click **Rename** to save the changes.

To change the Fabric Name, right-click the Fabric and select **Properties**. You can change the Name, Description, Location and Contact by double-clicking the desired field. Click **OK** after the changes are done.

## Changing the view for FICON management

To change the view for FICON, click **SAN** → **Options**. The Option window appears. Select **Display** and click the **Set Up FICON Display** button. Click **OK** to close the window. As shown in Figure 8-16 on page 261

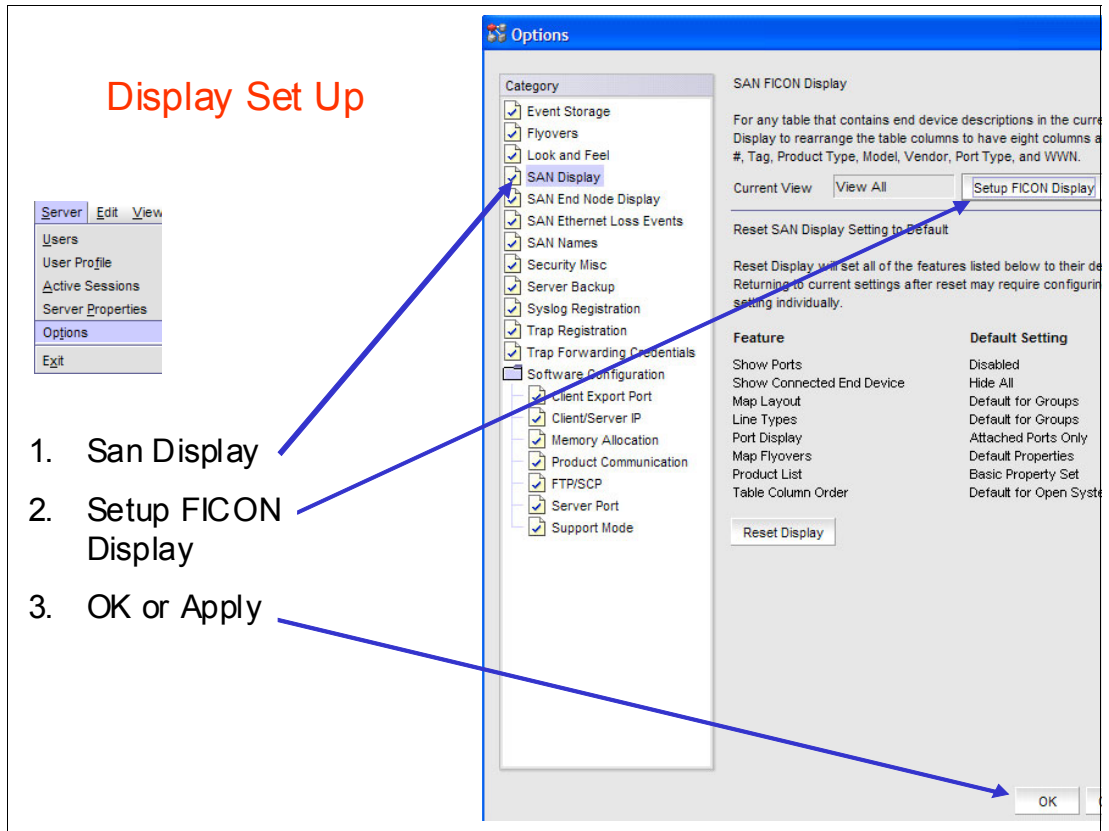


Figure 8-16 Ficon Display Setup

### Setting the display to hex

To view the display in hexadecimal, select the hex display option from the pull-down menu on the toolbar as illustrated in Figure 8-17.




Figure 8-17 Set to hex

### Viewing node descriptors

If you have changed the view for FICON management, you will see the node descriptor details in the IBM Network Advisor for each port on the left side of the window. To change the Port view, right-click a Director (on the left of the window) and select **Port Display**. Check-mark the details you want to display.

You can also display the node descriptor for a Director, by right-clicking it and selecting **Element Manager** → **Hardware**. In the Directors Element Manager, click **Name Server** on the left side of the window. The Name Server window opens and you will see all connected devices. By selecting the list entry and clicking **Detail View**, you will get a detailed view for the selected device.

## 8.3 Setting up a FICON Director

In this section, we illustrate how to set up FICON Directors, which applies to Directors in both topologies. We also include optional steps for the most commonly used functions.

Additional configuration steps for a cascaded FICON Director are covered in 8.4, “Setting up a cascaded topology” on page 298.

### 8.3.1 Changing the IP addresses

Before connecting the FICON Director to the network, change the default IP address to match your network layout and rules.

The steps to configure the Directors IP addresses are:

- ▶ Connecting a mobile or personal computer to the Director’s active Control Processor (CP)
- ▶ Configuring the mobile or personal computer to be in the same network
- ▶ Altering the Director’s IP addresses
- ▶ Connecting the Director to the customer network

**Note:** If you are redeploying a FICON Director and do not know its IP addresses, the serial port is used.

Steps to configure the Directors IP addresses using the serial port are described in *IBM System Storage SAN768B-2 Installation, Service, and User's Guide*, GA32-0893 and in *IBM System Storage SAN384B-2 Installation, Service, and User's Guide*, GA32-0894.

To change the IP addresses in the Director:

1. Connect a personal or mobile computer to the active Control Processor Card (CP0 in Slot 6 or CP1 in Slot 7). Use the Management Port to connect the network cable. The active CP is indicated by the blue LED on the card.
2. Alter the computer’s IP address to be in the same network as the Director. Select **Start** → **Control Panel** → **Network Connections** (use the classic view in the control panel).
3. Right-click the Network Adapter, and click **Properties**, as shown in Figure 8-18 on page 263.

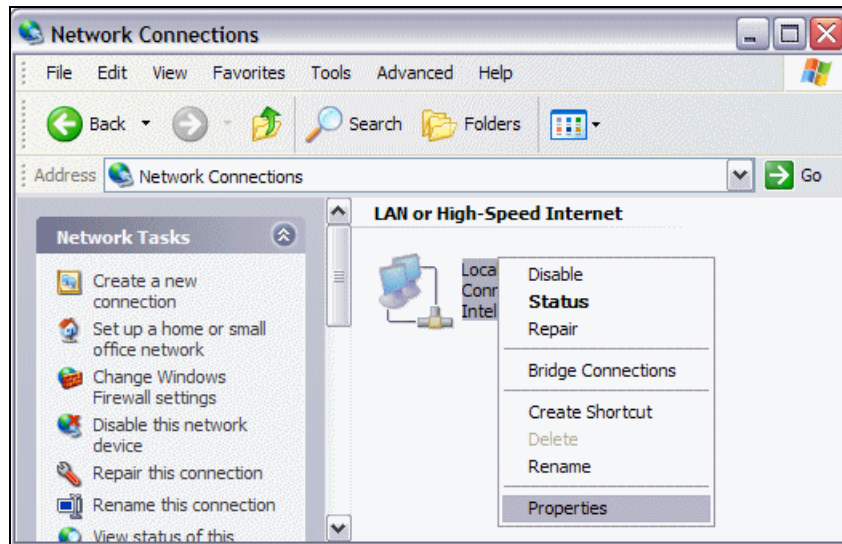


Figure 8-18 Select network interface

4. Select **Internet Protocol (TCP/IP)**, and click **Properties**, shown in Figure 8-19.

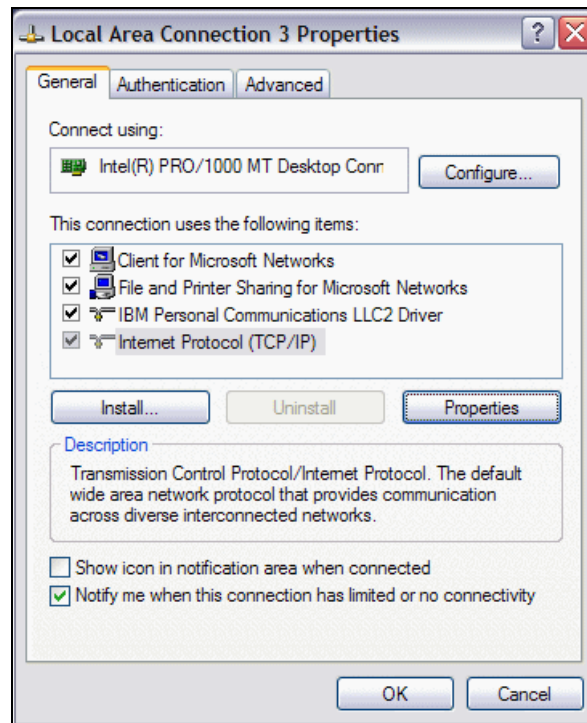


Figure 8-19 Select Protocol

5. Select **Use the following IP address**, and enter the address (172.126.20.81, in this case) to be in the same network as the Director, as shown in Figure 8-20 on page 264.
6. Click **OK** for both Properties windows.




Figure 8-20 Setting up a mobile computer IP address

7. Connect the cable to the active CP Card in the Director and open a web browser. Enter the default IP address in the web browser's address field (9.12.4.136, in this case), and press Enter.
8. In the login window, enter the default user (admin) and the default password (password), and then click **OK**.

(The Mozilla Firefox web browser asks for open or save. Select **open with** and click **OK**.)

If you receive a message like the one shown in Figure 8-21, you must install a new Java version (Version 1.6 or later). If you do not receive this message, proceed to step 9 on page 265.




Figure 8-21 Wrong Java version message

If Version 1.6 is installed, select **Start** → **Control Panel** → **Folder Options** (in classic view). Select the **File Types** tab, scroll down, and select **JNLP**. Click **Advanced**. In the Edit File Type window, select **Launch**, and click **Edit**. The window shown in Figure 8-22 on page 265 will display.

In the field Application used to perform action, you must change the location of your Java version, for example, C:\Program Files\Java\jre1.6.0\_05\bin\javaws.exe" "%1.



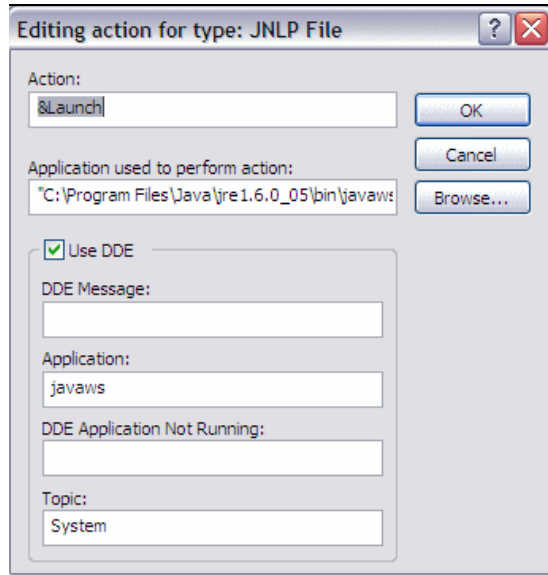


Figure 8-22 Assign Java version to JNLP files

Click **OK** after the changes are made.

The login window will display after you enter the IP address in the web browser's address field. Type in the user (admin) and password (password), and click **OK**.

9. At the Director's graphical user interface, select **Manage** → **Switch Admin**. Select the **Network** tab at the Switch Administration window to reach Figure 8-23 on page 266.

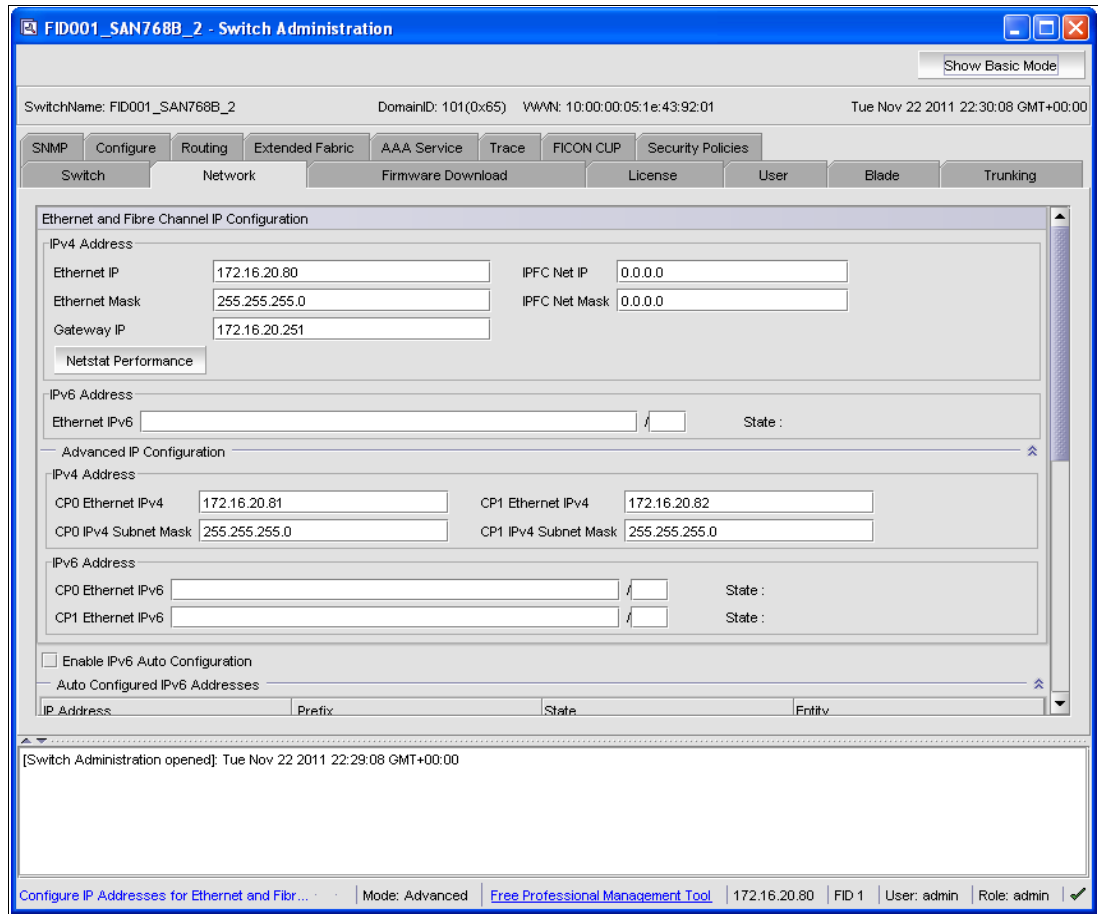


Figure 8-23 Setup management IP address

10. Populate the required fields.

Our example is based on Figure 8-2 on page 248. Table 8-1 lists the IP addressing scheme we used.

Table 8-1 IP address example

|                      | Chassis address | CP0 address   | CP1 address   |
|----------------------|-----------------|---------------|---------------|
| <b>Ethernet IP</b>   | 172.16.20.80    | 172.16.20.81  | 172.16.20.82  |
| <b>Ethernet mask</b> | 255.255.255.0   | 255.255.255.0 | 255.255.255.0 |
| <b>Gateway IP</b>    | 10.1.1.1        | n/a           | n/a           |

**The Chassis address:** The Chassis address is later used to connect to the Director to guarantee a connection even if a Control Processor switchover occurred (for example, during a firmware upgrade). When connecting with the IP address of CP0 or CP1, only the active CP is manageable.

IP addresses between 10.0.0.0 to 10.0.0.255 cannot be used because they are already used internally in the Director.

11. To activate the new IP addresses, scroll down to the bottom of the window, and click **Apply**.
12. Connect each Control Processor Management Port to the network (Ethernet switch). The Director needs to be discovered by IBM Network Advisor, as described in “Discovering a FICON Director” on page 259.

### 8.3.2 Using the FICON wizard

The FICON wizard provides you with a simple means to set up the following configuration options:

- ▶ Insistent Domain ID
- ▶ Set high integrity fabric
- ▶ Enable port-based routing
- ▶ Enable in order delivery
- ▶ Disable dynamic load sharing
- ▶ Enable the FICON Management Server (FMS) for Control Unit Port (CUP) management

The configuration wizard can be used to make these listed settings for the configuration and ensures that Lossless is configured with DLS so there are fewer opportunities for mistakes. Furthermore, the wizard makes these changes on all switches in the fabric, eliminating the need to manually make these changes again for other switches if cascaded and ensures all settings are the same on all switches in the fabric.

You can access the FICON wizard using IBM Network Advisor by right-clicking the Director and selecting **Configure** → **FICON** → **Configure Fabric**, as shown in Figure 8-24.

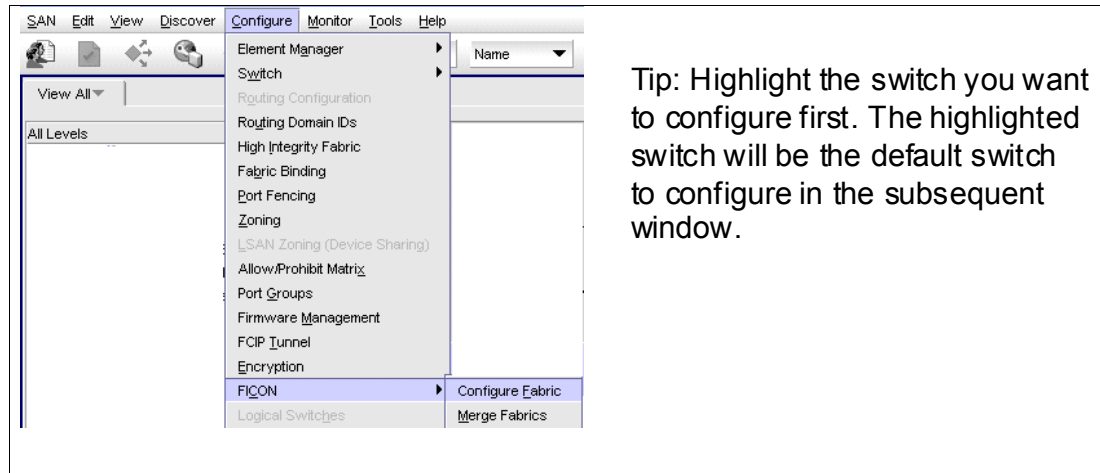


Figure 8-24 Configure Fabric wizard option

Figure 8-25 on page 268 is an example of the setting, which can be enabled using the Configure Fabric Wizard.

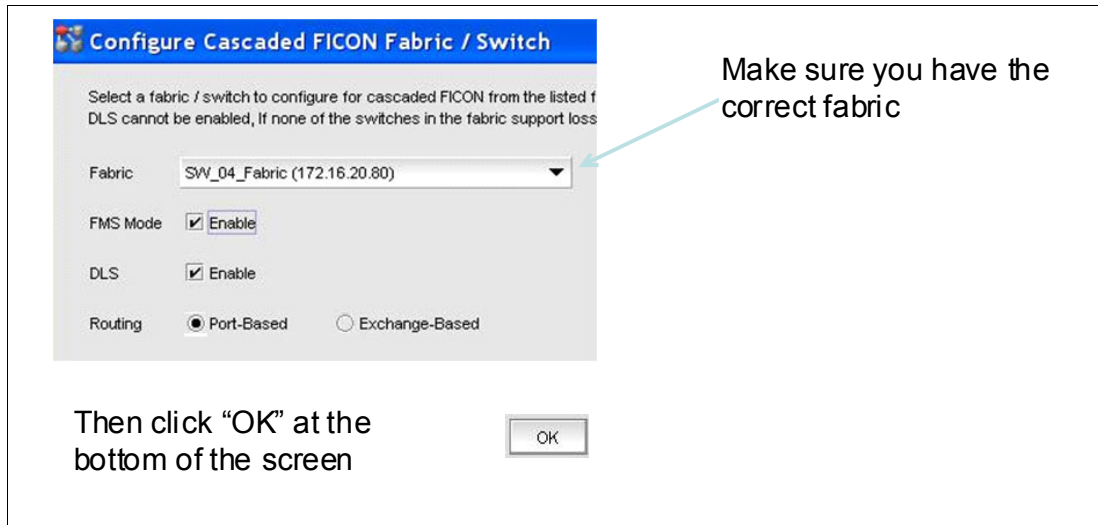


Figure 8-25 Configure Fabric wizard panel

**Note:** With FOS 7.0 and above, FICON Management Server (FMS) mode cannot be set without the applicable license, so to enable it the license must be installed first. See 8.3.3, “Enabling features (optional)” on page 268.

If you decide not to use the FICON wizard or you need to configure other options not included in the FICON wizard, refer to the remaining parts of this section.

### 8.3.3 Enabling features (optional)

To use a licensed feature, it must first be enabled:

1. Go to the Switch Admin window in the Element Manager of the Director. This can be done using IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. Select the License tab, as shown in Figure 8-26 on page 269.

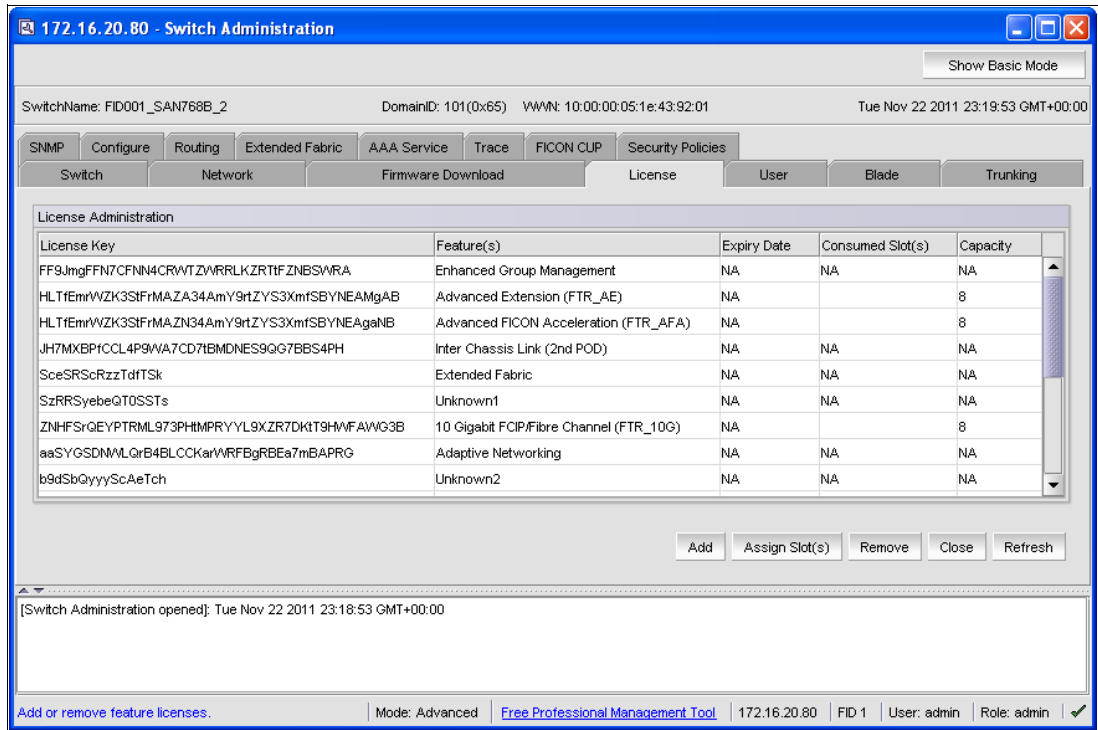


Figure 8-26 License overview

- Click **Add**, and then enter the license key for the feature you want to enable, as shown in Figure 8-27.

**Note:** The license key is case sensitive and will only work on the Director to which it is assigned. Therefore, you must compare the WWN on the license with your Director.

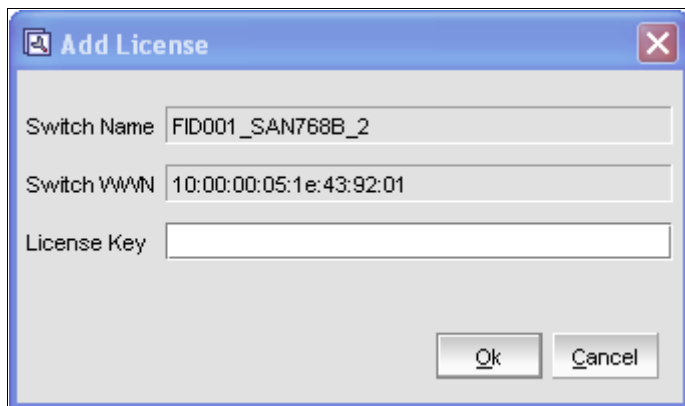


Figure 8-27 Add a license

- Click **OK** to activate the license. If you have more than one license, repeat step 3 for each. The feature is now enabled and can be used or configured.

### 8.3.4 Setting up a logical switch (optional)

A logical switch in the Director can be used to separate the physical ports for different System z environments or for mixed environments (FICON and Fibre Channel).

After Virtual Fabric (VF) is enabled, up to seven additional Logical Switches (LS) can be created (switch model dependent). Two of these switches can have FICON Management Server and CUP enabled. The default logical switch has Logical Fabric ID (FID) 128 assigned. If you want to assign port IDs (PIDs), you can configure a logical switch and assign all ports from the default logical switch to the new logical switch. The default logical switch is always present if Virtual Fabric is enabled, but it does not need to have ports assigned to it.

If the Director uses 48 port cards, you must create two logical switches with Zero-based Area Assignment. If the Director was *delivered* with Firmware level 6.2.0 or above (required for logical switches), you have by default one logical switch, thereby allowing you to create more.

**Important:** If the Firmware was *updated* to level 6.2.0, then you need to enable Virtual Fabric first, which is disruptive.

To create a logical switch:

1. If Virtual Fabric is enabled (indicated by the blue V on the Director in the IBM Network Advisor Switch view), proceed to step 6 on page 272; otherwise, select the Director in IBM Network Advisor that you want to enable for Virtual Fabric in the Chassis Group.
2. Select the Director, and select **Configure** → **Virtual Fabric** → **Enable** at the menu bar, as shown in Figure 8-28.

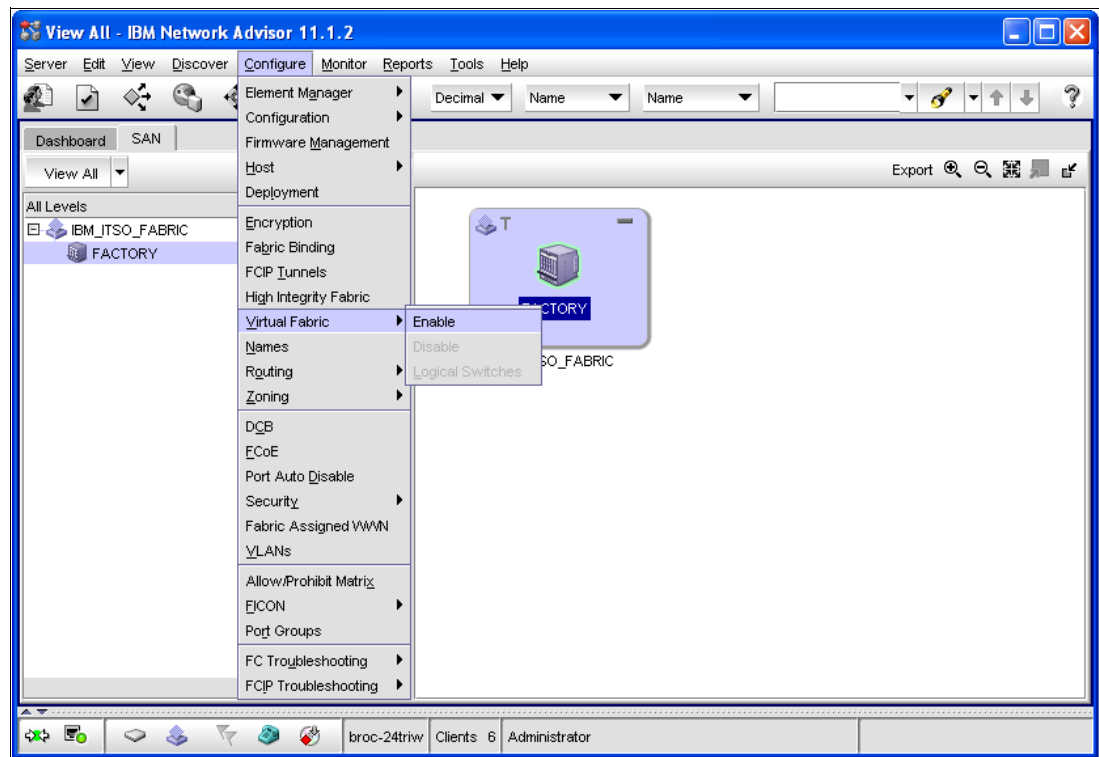


Figure 8-28 Enable Virtual Fabric

3. Read the warning message in Figure 8-29 on page 271, and select **OK**.



Figure 8-29 Enable VF Warning message

**Important:** Clicking **OK** immediately reboots the Director and creates a logical switch, which is disruptive to the Director.

4. Wait for the Director to complete the reboot. You must discover the Director again after the reboot because the Director was changed and appears as a new Director. Refer to “Discovering a FICON Director” on page 259 for more details about this topic.
5. The Director shows that the Virtual Fabric is enabled by a blue V in the IBM Network Advisor switch view, as shown in Figure 8-30.

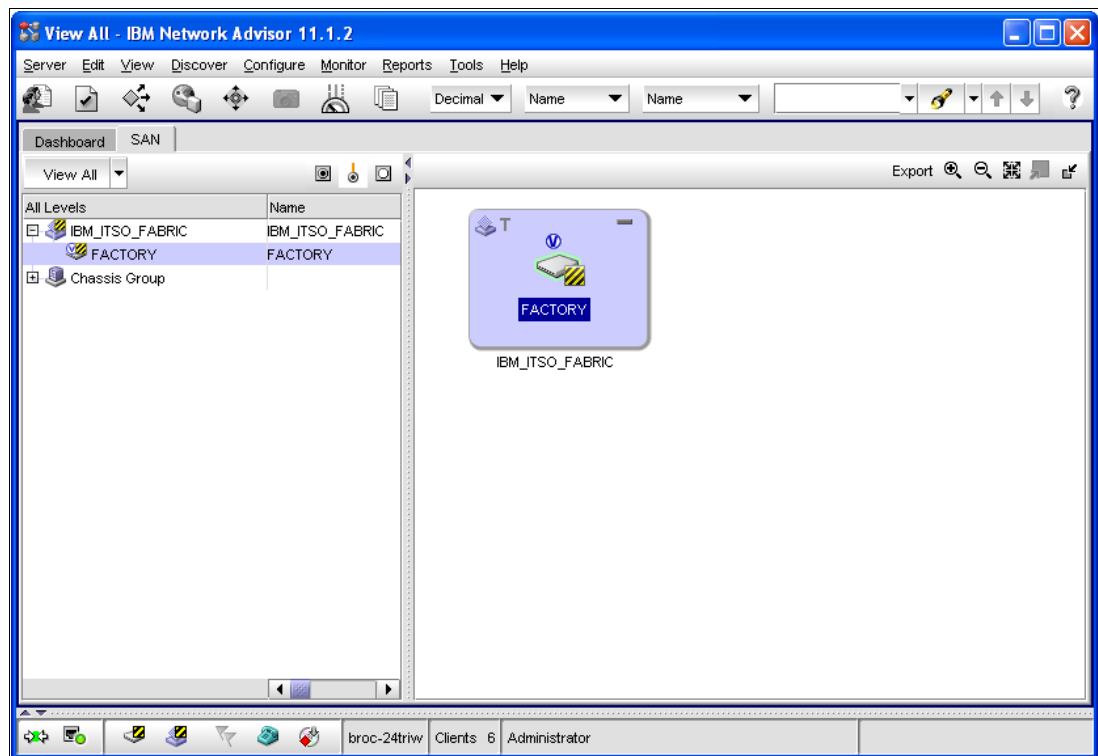


Figure 8-30 Virtual Fabric enabled

6. Select **Configure** → **Logical Switches** at the IBM Network Advisor menu bar.
7. Wait for the Logical Switches configuration window to appear, as shown in Figure 8-31.
8. In the Chassis drop-down menu, select the chassis where you want to create a new logical switch.

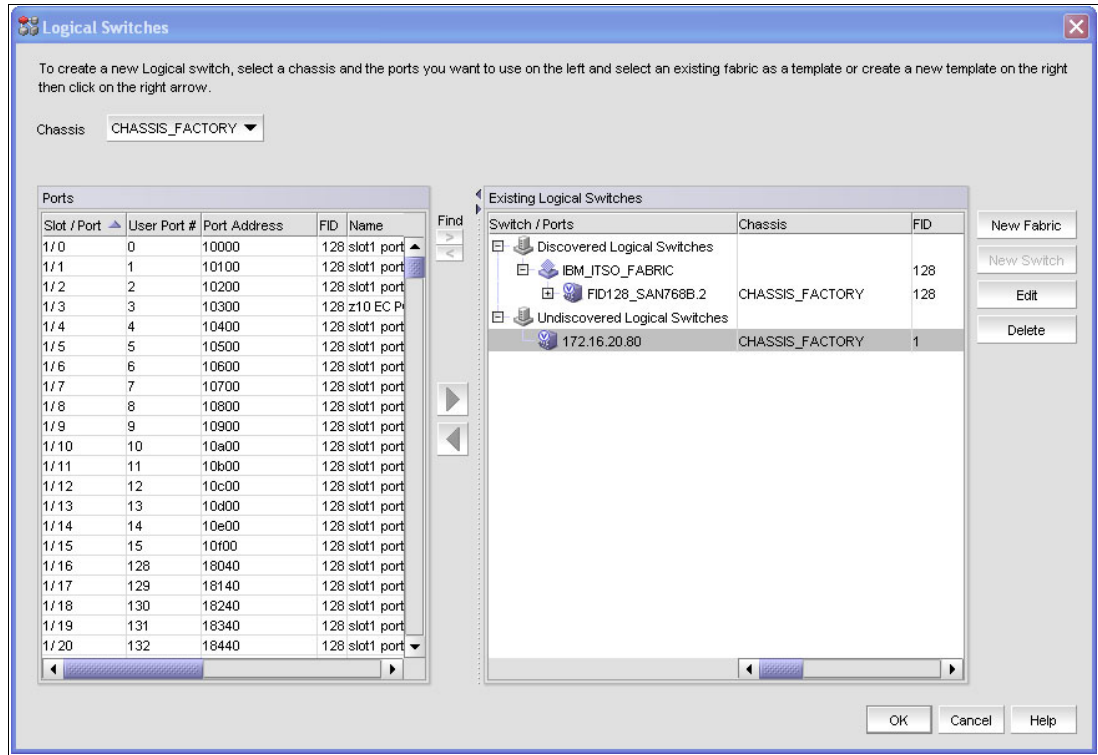


Figure 8-31 Create a Logical Switch

9. Select **Undiscovered Logical Switch**, and then click **New Switch**.
10. Enter the desired settings for the fabric using the Fabric tab, as shown in Figure 8-32.

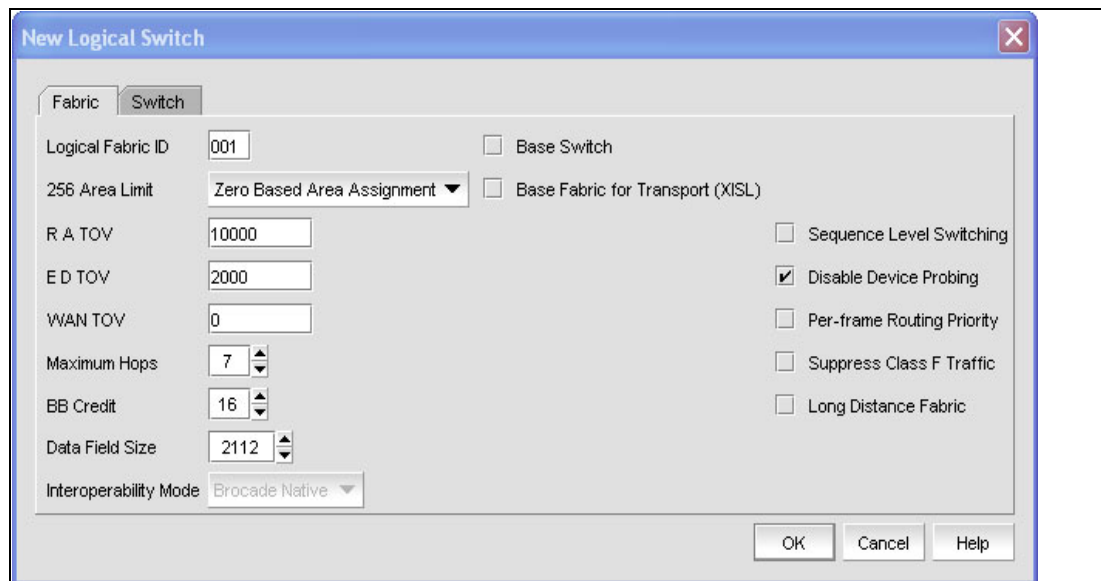


Figure 8-32 Logical switch configuration



**Note:** R A TOV, E D TOV, WAN TOV, Maximum Hops, BB Credit, and Data Field Size should be left at the defaults illustrated here unless otherwise advised by an IBM or Brocade-certified fabric professional. Increasing BB credits for ISL links is performed in another step and must not be changed here.

The Logical Fabric ID (FID) can be a number from one to 128 (decimal). The FID must be unique in the selected chassis.

**Note:** Keep in mind that only logical switches with the same FID can be cascaded.

Do not select Base Switch or Base Fabric for Transport, and do not set the 256 Area Limit to Disable, because these are not supported by FICON.

The 256 Area Limit can be set to Zero-Based Area Assignment. Ports are numbered in ascending order, starting from 00 to FF, for each logical switch.

You can change the port numbering for logical switches (except for the default logical switch) later in the Port Admin panel in the Element Manager by using the Bind PID and Un-Bind PID buttons.

11. Select the **Switch** tab and assign the Domain ID and the logical switch name. Also, check-mark the Insistent Domain ID field.
12. Click **OK** after all settings are complete.
13. Assign the desired ports to the new logical switch. Select the ports on the left side, and click the arrow pointing to the right, as shown in Figure 8-33.

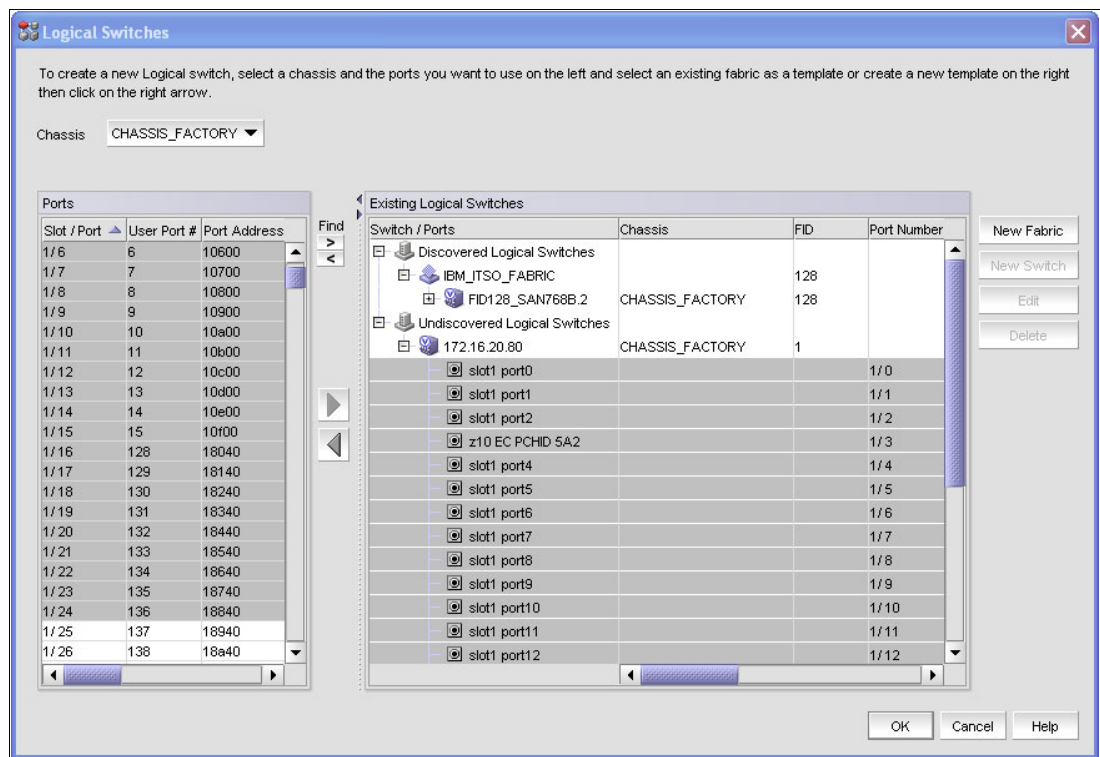


Figure 8-33 Assign Logical ports

14. After the ports are assigned, click **OK**.
15. Verify the changes in the confirmation window, and select **Re-Enable Ports** after moving them (if not selected, all ports must be enabled manually). Click **Start**. See Figure 8-34.

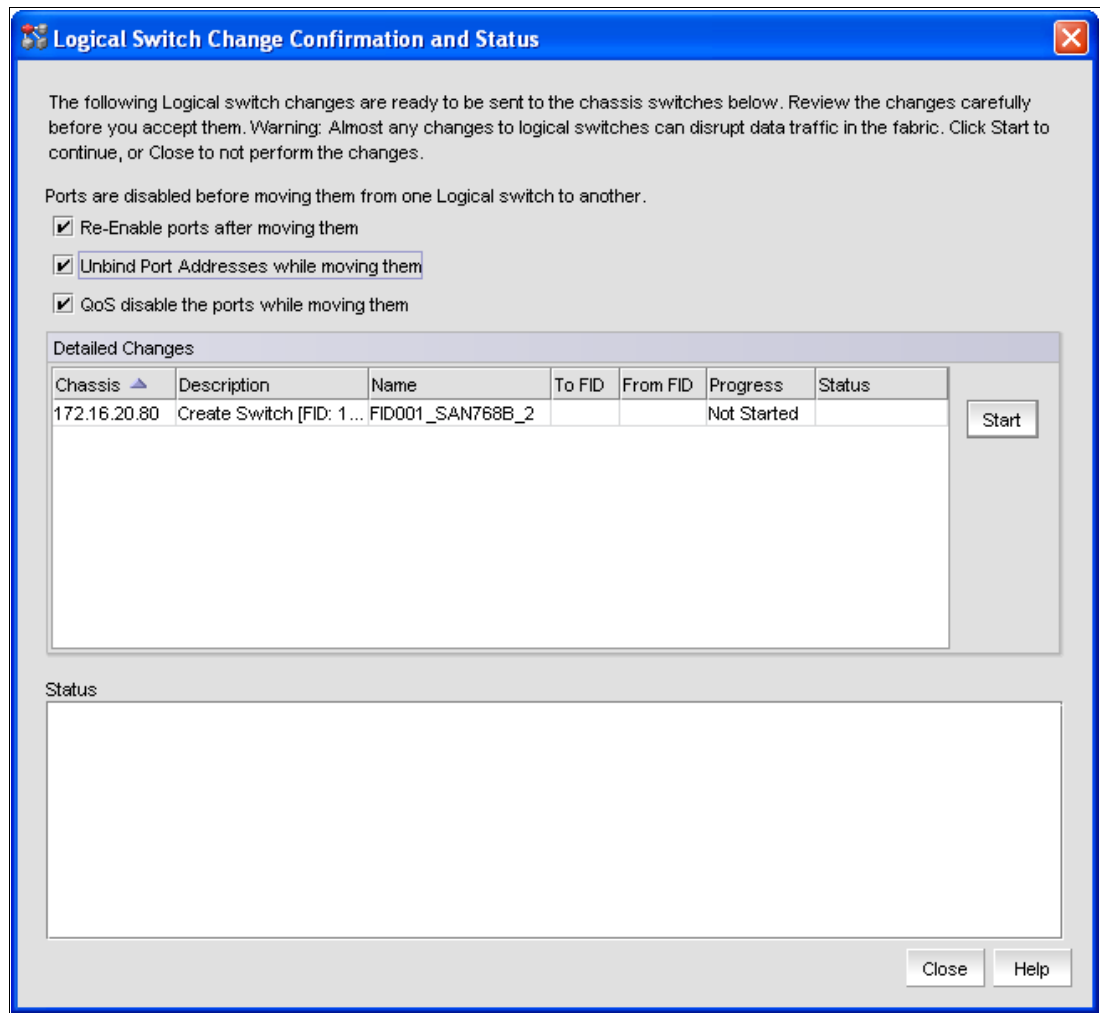


Figure 8-34 Logical Switch Change Confirmation and Status

16. Wait for the changes to complete. After the status Success displays, click **Close**. See Figure 8-35 on page 275.

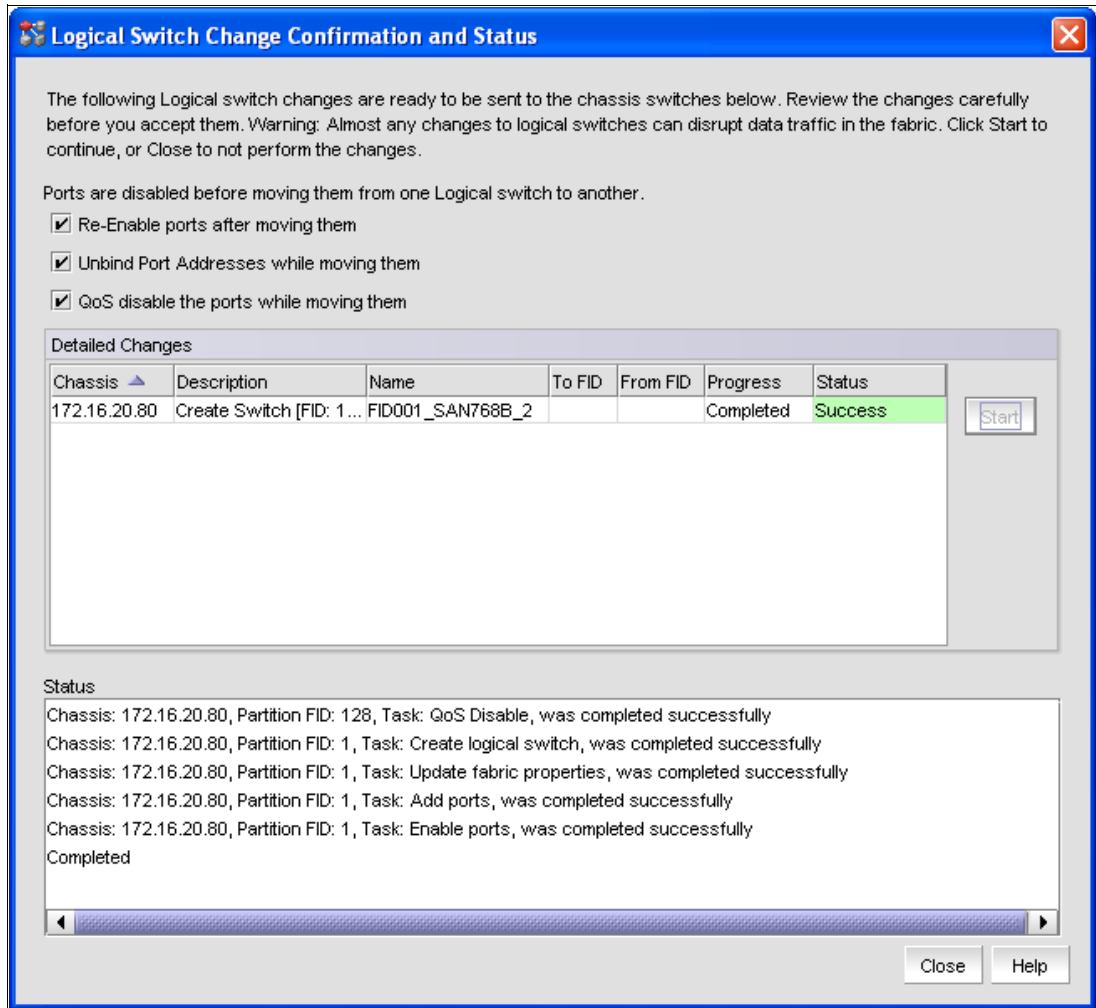


Figure 8-35 Logical Switch Change Confirmation and Success State

17. You must discover the new logical switch. Instructions that explain this process are provided in “Discovering a FICON Director” on page 259.

The logical switch is now created and can be used like any other physical switch or Director.

### 8.3.5 Configuring the Domain ID and Insistent Domain ID

To set up the Domain ID, the Director must be disabled, as explained here. We also explain how to set the Insistent Domain ID flag, which allows you to later merge two Directors without disruption. Only Domain IDs between one and EF in hex are possible, which is decimal one to 239. The default Domain ID is one.

To change the Domain ID to the desired value:

1. Go to the Switch Admin window in the Element Manager of the Director. This can be done through the IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. Select the **Switch** tab, and then select **Disable** in the Switch Status field. lick **Apply**, as shown in Figure 8-36 on page 276.

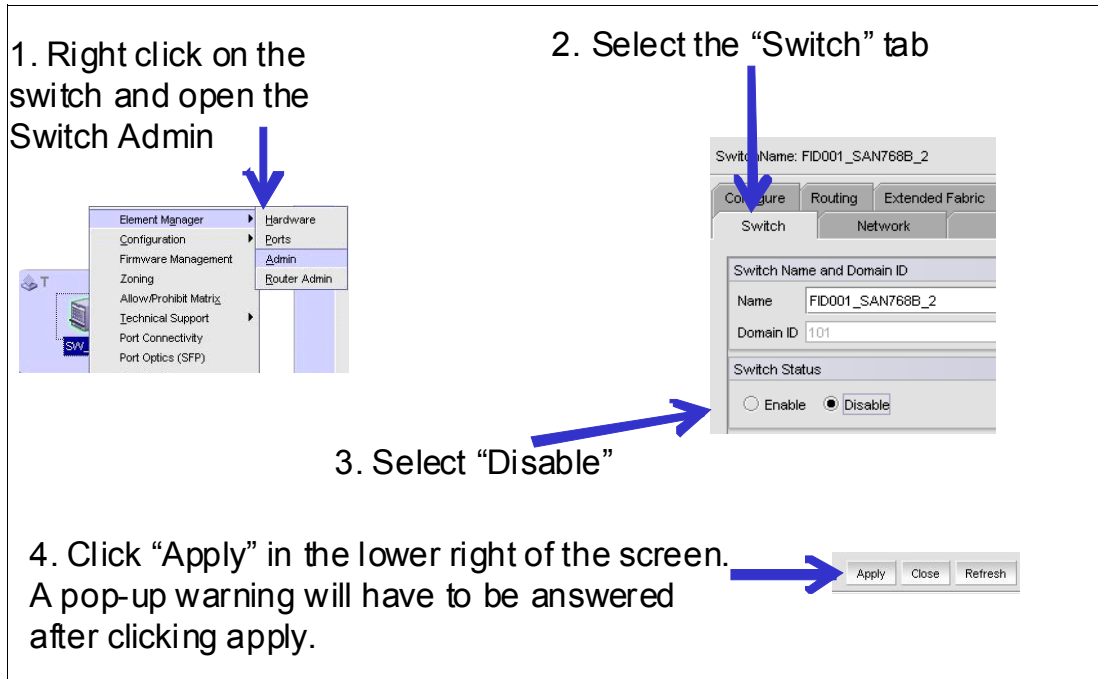


Figure 8-36 Domain ID configuration

**Important:** Setting the Director to disable is disruptive.

3. Read the warning message, and confirm it by clicking **Yes**.

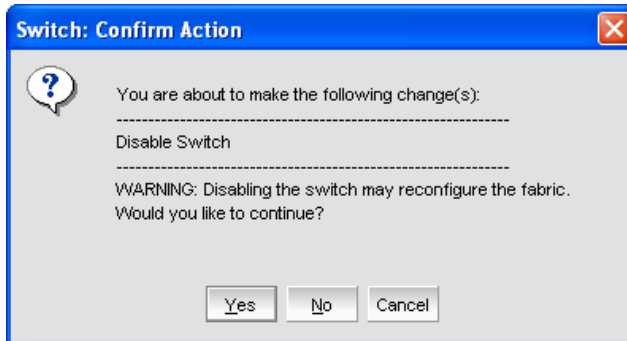


Figure 8-37 Disable switch warning

4. Change the Domain ID field to the desired value. In our example, it is 65 in hex, which is 101 in decimal. The value must be entered in decimal. You can also change the Name of the Director so it can be identified easily.

**Note:** The Domain ID must match the switch address configured in HCD, HMC or IOCP on the System z side. Be aware that the switch address on the System z side is configured in hex, but the Domain ID in the Director is defined with a decimal value.

5. Click **Apply**. After reading the warning message, click **Yes** to confirm the action.
6. If not already done, click **Show Advanced Mode** in the upper-right corner of the Switch Administration window.

7. Click the **Configure** tab.
8. Click **Insistent Domain ID Mode** to check-mark it, as shown in Figure 8-38.

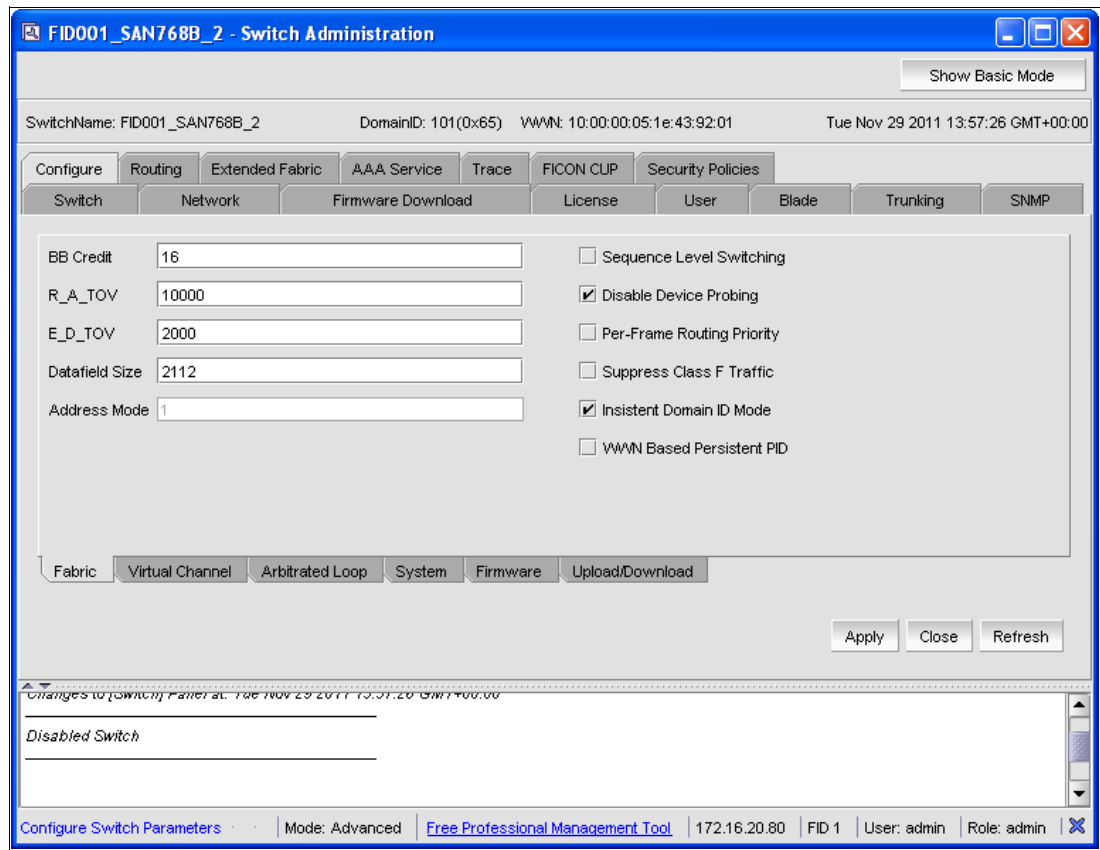


Figure 8-38 Insistent Domain ID Mode

**Note:** Do not change BB credits here. Increasing BB credits for ISLs will be addressed later. All other parameters should be as shown in Figure 8-38. Making other changes here is rare and should only be done so with the advice of IBM support.

9. Click **Apply** to save the changes.
10. To enable the Director at the Switch Status field, return to the Switch tab. (You can also skip this step and configure Port-Based-Routing (PBR), Dynamic Load Sharing (DLS) and In-Order Delivery (IOD), before you enable the Director.)
11. Select **Enable** in the Switch Status field, and click **Apply**.
12. Read the warning message, and click **Yes** to confirm the action.

The Domain ID is now set. It cannot be automatically changed if another Director joins the Fabric with the same Domain ID; instead, the new Director is segmented until another Domain ID is set at the joining Director.

### 8.3.6 Setting up PBR, IOD, and DLS

These settings can be configured using the FICON wizard, as detailed in 8.3.2, “Using the FICON wizard” on page 267.

Port-Based Routing (PBR) and In-Order Delivery (IOD) must be set in a FICON environment. Dynamic Load Sharing (DLS) can be used in a FICON environment if lossless DLS is also enabled; otherwise, it must be set to disabled.

Lossless DLS ensures that no frames are dropped during rebalancing of the Inter-Switch Links (ISLs), which is the case if an ISL is added or removed. However, lossless DLS is not supported in long distance environments. When lossless DLS is used in long distance environments, you must disable credit recovery at all ISLs.

To set up PBR, IOD, and DLS:

1. Return to the Switch Admin window in the Element Manager of the Director. This can be done using the IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. If the Director is already disabled, go to step 6. Otherwise, select the **Switch** tab.
3. Select **Disable** in the Switch Status field, and click **Apply**.

**Important:** Setting the Director to disable is disruptive.

4. Read the warning message, and confirm it by clicking **Yes**.
5. Click **Show Advanced Mode** in the upper-right corner.
6. Select the **Routing** tab, as shown in Figure 8-39.

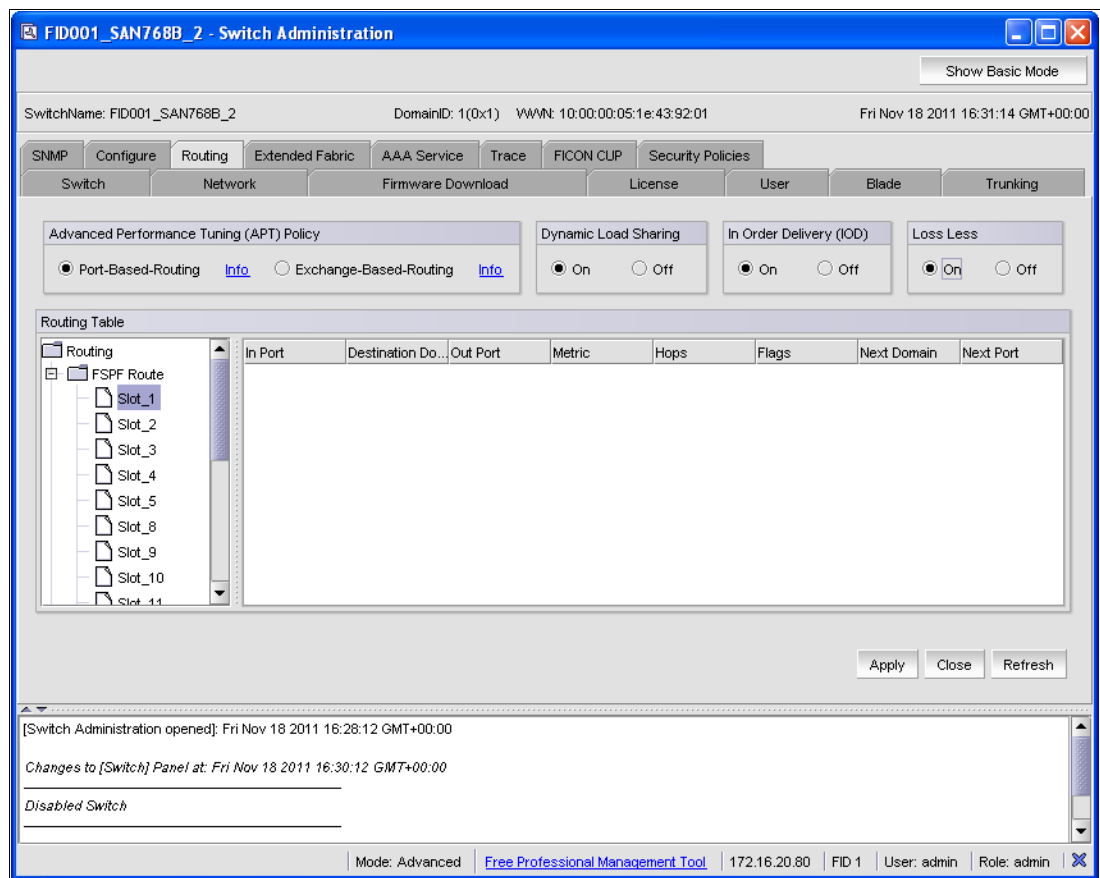


Figure 8-39 Setup Port-based Routing, DLS, and IOD

7. Select **Port-Based-Routing** in the Advanced Performance Tuning (APT) Policy field. Set Dynamic Load Sharing (DLS) to **Off** if lossless DLS must not be used. Set In-Order Delivery (IOD) to **On**. Click **Apply** to save the changes.
8. If you set DLS to **On**, Set In-Order Delivery (IOD) and Lossless DLS to **On**. Click **Apply** to save the changes.
9. For long distance environments, execute `portcfgcredirecovery --disable slot/port` for each ISL, using the command-line interface (CLI). Otherwise, proceed to step 11.
 

The following steps explain one method of using the CLI:

  - a. In the IBM Network Advisor server or client, right-click the Director, and select **Telnet**. A new window opens. (You can also use an `ssh` connection to the Director, if needed.)
  - b. Enter the user name (in our case, `admin`), and press **Enter**.
  - c. Enter the password, and press **Enter** (the default is password).
  - d. Enter, for example, `setcontext 1` to connect to a logical switch with FID 1.
  - e. For long distance environments, execute `portcfgcredirecovery --disable slot/port` for each ISL.
  - f. Type `logout`, and close the window.
10. To enable the Director at the Switch Status field, return to the **Switch** tab. (You can also skip this step and configure the Control Unit Port (CUP) before enabling the Director.)
11. Select **Enable** in the Switch Status field, and click **Apply**.
12. Read the warning message, as shown in Figure 8-40, and click **Yes** to confirm the action.

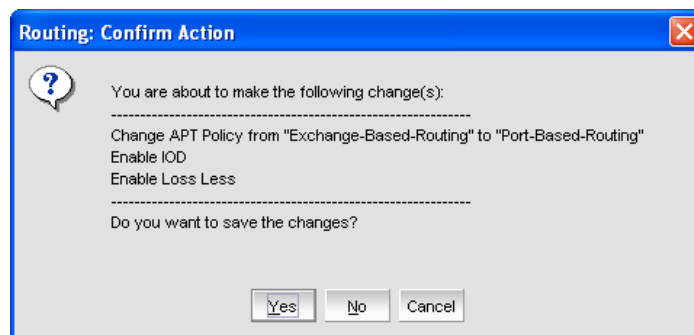


Figure 8-40 Confirm changes

The Director is now ready to be used in a FICON environment.

### 8.3.7 Enabling the Control Unit Port

**Note:** With FOS 7.0 and above, FICON Management Server (FMS) mode cannot be set without the applicable license, so to enable it the license must be installed first. See 8.3.3, “Enabling features (optional)” on page 268.

FMS can be configured using the FICON wizard, as detailed in 8.3.2, “Using the FICON wizard” on page 267.

If the Director has a CUP license, the Control Unit Port (CUP) can be enabled. In this case, ports FE and FF in a 256 Port Director must be disabled because those port addresses are

used by the CUP server. Refer to 8.3.8, “Changing the port type and speed” on page 282 for more information.

The Director must be enabled before you can disable the ports. The CUP can be enabled while the Director is enabled. A Registered State Change Notification (RSCN) is sent to inform the System z server that the CUP is now enabled. However, because our example is in the process of setting up the Director, you can do it while the Director is disabled.

To enable the CUP:

1. Go to the Switch Admin window in the Element Manager of the Director using the IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. If the Director is already disabled, proceed to step 6. Otherwise, select the **Switch** tab.
3. Select **Disable** in the Switch Status field, and click **Apply**.

**Important:** Setting the Director to disable is disruptive.

4. Read the warning message, and confirm it by clicking **Yes**.
5. Click **Show Advanced Mode** in the upper-right corner.
6. Select the **FICON CUP** tab, as shown in Figure 8-41.

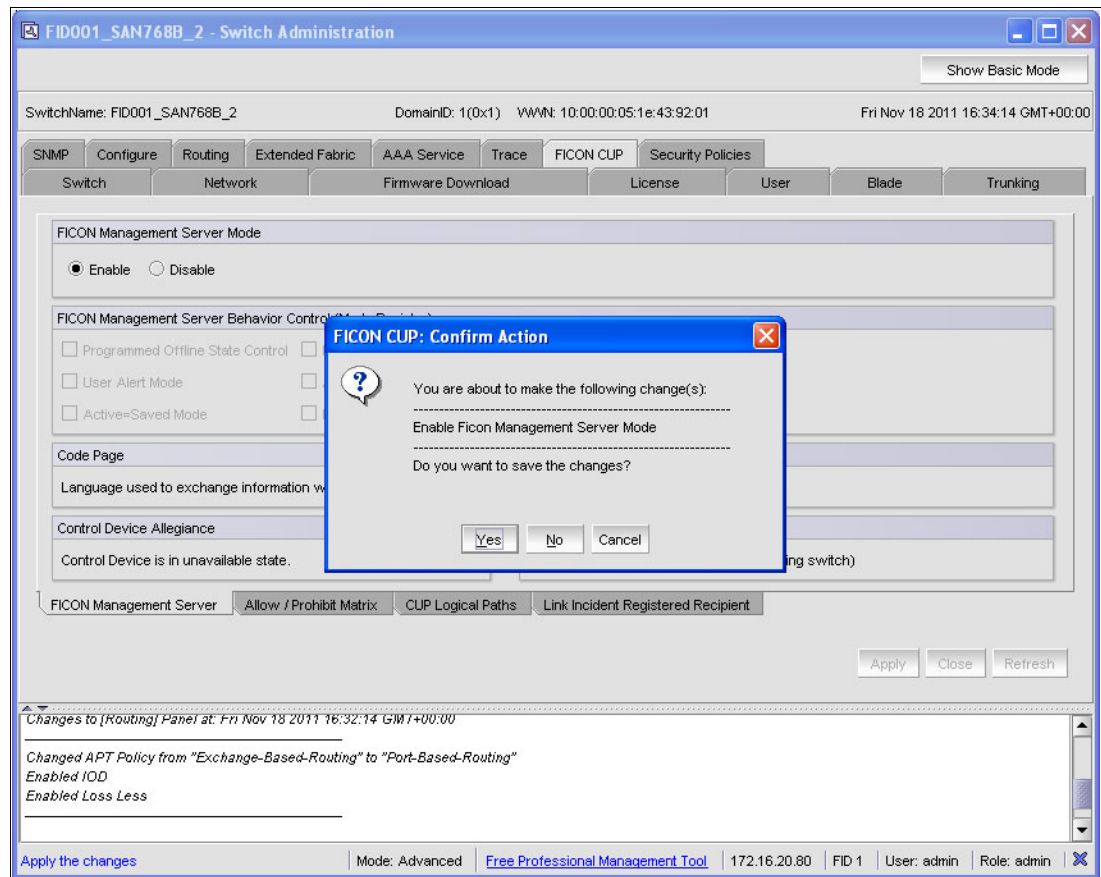


Figure 8-41 Enable the CUP

7. Select **Enable** in the FICON Management Server Mode field.



8. Click **Apply**.
9. Read the message, and click **Yes** to confirm the action.
10. Now you can alter the FICON Management Server Behavior Control (Mode Register) field. Select the desired functions you want to use with the CUP, as shown in Figure 8-42.




Figure 8-42 FICON Management Server Behavior Control

11. Click **Apply** after the changes are made.
12. Read the message, and click **Yes** to confirm the action.
13. Return to the **Switch** tab to enable the Director at the Switch Status field.
14. Select **Enable** in the Switch Status field, and click **Apply**.
15. Read the warning message, and click **Yes** to confirm the action.

The CUP Port is now usable from the System z server.

### Missing Interrupt Handler Process Timeout (MIHPTO)

The default timeout value for the CUP is 180 seconds. To check the value, or to change it, you must connect using telnet or ssh to the Director's IP address.

Here is one method for accomplishing this task:

1. In the IBM Network Advisor server or client, right-click the Director, and select **Telnet**. A new window opens. (You can also use an ssh connection to the Director, if needed.)
2. Enter the user name (in this case, admin), and then press **Enter**.
3. Enter the password, and press **Enter** (the default is password), as shown in Figure 8-43.

4. Enter, for example, **setcontext 1**, if you need to connect to a logical switch with FID 1.

```
FID001_SAN768B_2:FID1:admin> ficoncupset mihpto 181

MIHPTO has been changed to 180 seconds

FID001_SAN768B_2:FID1:admin> ficoncupshow mihpto

MIHPTO for the CUP: 180 seconds

FID001_SAN768B_2:FID1:admin> ficoncupset mihpto 190

MIHPTO has been changed to 190 seconds

FID001_SAN768B_2:FID1:admin>
```

Figure 8-43 FICON CUP MIHPTO

5. To see the timeout that is currently set, enter **ficoncupshow mihpto**.
6. To change the value, enter **ficoncupset mihpto xxx**, where xxx is the new value. This value must be divisible by 10 if over 63, or will be rounded down. A value between 15 and 600 seconds can be used.
7. Enter **logout**, and close the window.

The MIHPTO for the CUP is now set and active.

### 8.3.8 Changing the port type and speed

To change a specific port type, speed, and BB credits you must modify the port configuration. The default port type is a U\_Port, which automatically detects if an F\_Port, E\_Port, or L\_Port device is connected. The default speed is set to auto-negotiate.

**Important:** Changing the port configuration is disruptive to port operations.

The recommended best practice is to leave the port type as a U\_Port and leave the speed at auto-negotiate. Some customers might want to set the port type as an added security measure. Setting the speed is typically only done in the rare instance when a device has difficulty auto negotiating or to limit network bandwidth. Typically, the only reason to use the port configuration wizard is to change the BB credits for ISL ports.

To configure a port:

1. Go to the port configuration window of the Director's Element Manager. Right-click the Director in the IBM Network Advisor, and select **Element Manager** → **Ports**. The Port Administration window is displayed.
2. Click **Show Advanced Mode** in the upper-right corner to see the configuration options.
3. Select the port you want to configure, and click **Disable**.

**Important:** This is disruptive for the traffic on the selected port.

4. Read the information at the confirmation panel, and click **Yes**.
5. Click **Edit Configuration** for the selected port, as shown in Figure 8-44.

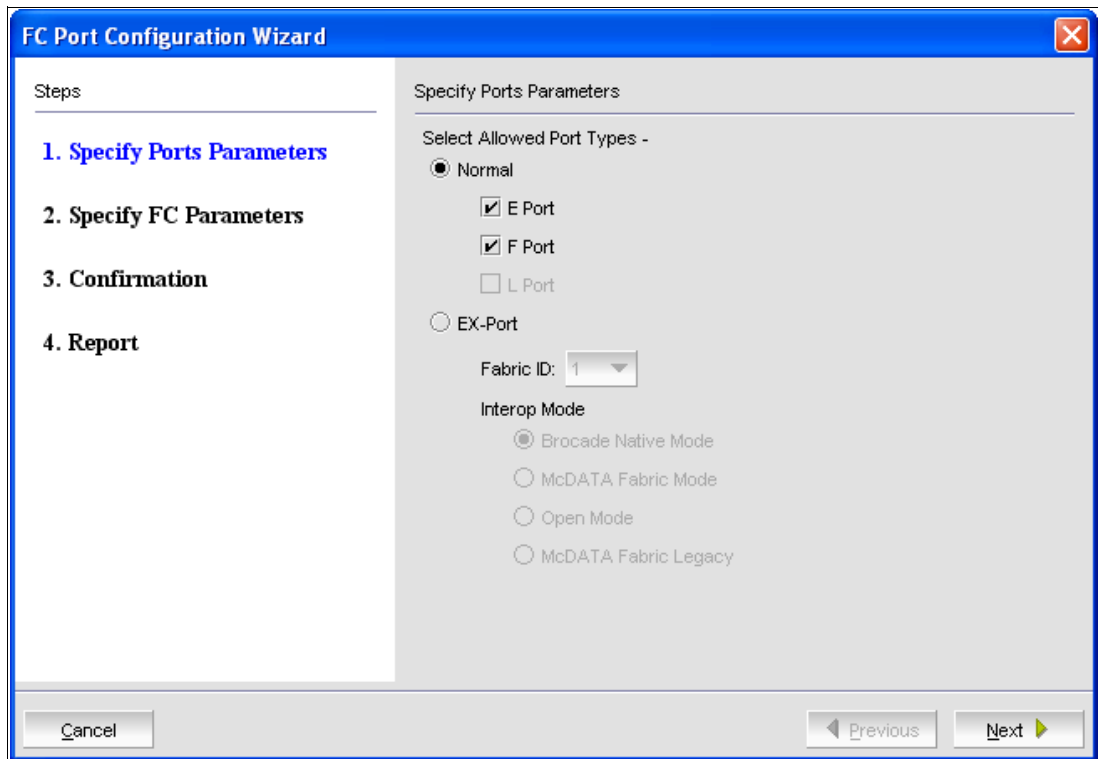


Figure 8-44 Port type configuration

6. Select the Port Type you want to configure, or accept the defaults if you only want to change the speed. Note the following points:
  - E\_Port is for ISLs.
  - L\_Port is for loop devices.
  - F\_Port is for all other devices.

Be aware that EX\_Port is *not* selectable because it is only available at the Base Switch in a logical switch configuration, which is not supported for a FICON environment.

7. Click **Next**, and Figure 8-45 on page 284 is displayed.

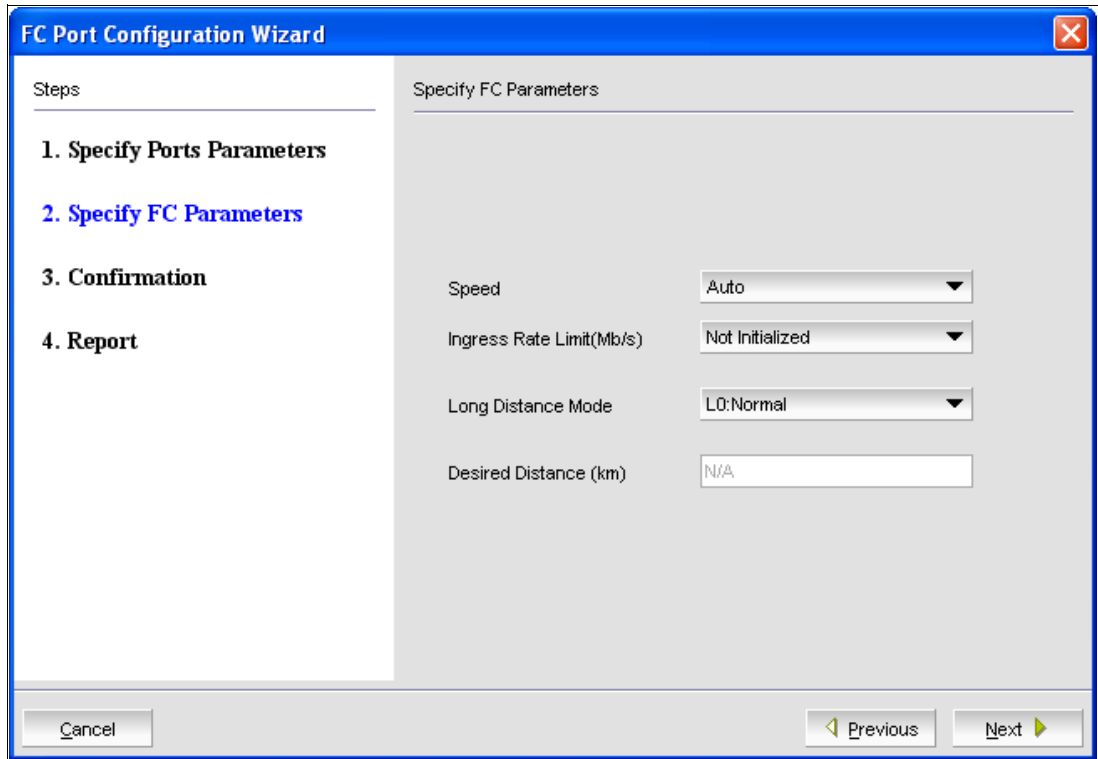


Figure 8-45 Port speed and distance configuration

Note the following points:

- At the Speed drop-down menu, you can select Auto (for auto-negotiate) or 2, 4, 8, and 16 Gbps.
- The Ingress Speed Limit (MBps) drop-down menu is a Quality of Service function, to limit the throughput of the port (200, 400, 600, 800, 1000, 1500, 2000, 3500, 4000, 5000, 7000, 8000, 9000, 10000, 11000, 12000, 13000, 14000, 15000, and 16000 Mbps are possible).
- With Long Distance Mode, you can configure the distance of the link (L0:normal is the default value). For more detailed information about this topic, refer to 8.3.9, “Changing buffer credits” on page 286.

8. Click **Next** after all selections are complete.
9. At the Confirmation window, verify the settings, and click **Save**.
10. Click **Close** after the configuration complete message.
11. Click **Enable** to enable the port after the configuration is complete.

After the port type and speed are configured, it is useful to give the port a name by clicking **Rename**.

### Fill word

A *fill word* is sent when there is no traffic on the link to keep it online. It is changed using the `portcfgfillword` command. However on the SAN768B-2 and SAN384B-2 switches, which are Condor-3 based platforms, this command is no longer supported because the fill word is set automatically. However, the following section provides details about changing fill words if relevant to your environment for SAN768B and SAN384B (Condor-2 based) FICON Directors.

For 8 Gbps SFPs, the fill word is ARB(FF), according to the specifications. However, some interfaces were developed before the final specifications were approved. These interfaces use *idles* to fill the link, which is the standard for 1 Gbps, 2 Gbps, and 4 Gbps.

FOS 6.2.0c introduced the **portcfgfillword** command, detailed in Table 8-2, which adds the ability to change the fillword setting. The fillword setting can be changed by enabling a particular numeric mode per port (0 or 1) or specifying the fillword value. From FOS 6.3.1 onwards, two additional fillword modes (2 and 3) were added, with mode 3 being the preferred mode that works with all 8 Gb devices. The mode can only be verified and changed through the CLI.

Table 8-2 *portcfgfillword mode parameter*

| Mode | Fillword value | Note   |
|------|----------------|--|
| 0    | -idle-idle     | Sets IDLE mode in the Link Init and IDLE as the fill word  |
| 1    | -arbff-arbff   | Sets ARB(ff) in the Link Init and ARB(ff) as the fill word   |
| 2    | -idlef-arbff   | Sets IDLE mode in the Link Init and ARB(ff) as the fill word   |
| 3    | -aa-then-ia    | Attempts hardware arbff-arbff (mode 1). If the attempt fails to go into active state, this command executes software idle-arb (mode 2) |

**Important:** Changing the mode is disruptive for the port.

The following method illustrates one way to change the fillword for a 8 Gbps interface:

1. In the DCFM server or client, right-click the Director and select **Telnet**. A new window opens. (You can also use an ssh connection to the Director, if needed.)
2. Enter the user name (in this case, admin), and press **Enter**.
3. Enter the password, and press **Enter** (the default is password).
4. Enter, for example, **setcontext 1** if you need to connect to a logical switch with FID 1.
5. The **portcfgfillword 1/8 0** command changes the fill word for Slot1, Port8 to Idles. The **portcfgfillword 1/8 3** command changes the fill word for Slot1, Port8 to mode 3.
6. The **portcfgshow** command will list all ports of the Director. The **portcfgshow 1/8** command will only list the settings for Slot1, Port8, as shown in Example 8-2.

Example 8-2 *Changing fill words*

```

-----
IBM_SAN768B:FID128:admin> portcfgfillword 1/8 3
IBM_SAN768B:FID128:admin> portcfgshow
Ports of Slot 1   0  1  2  3   4  5  6  7   8  9 10 11  12 13 14 15
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Speed           AN AN AN AN  AN AN AN AN  AN AN AN AN  AN AN AN AN
Fill Word       0  0  0  0   0  0  0  0   3  0  0  0   0  0  0  0

IBM_SAN768B:FID128:admin>logout

```

7. You can also use a script that changes all ports on a given port card, as shown in Example 8-3 on page 286.

*Example 8-3 Script to change the fill word for a 16-port card in slot 2*

---

```
for ((i=0;i<16;i++)); do (echo 2/$i;portcfgfillword 2/$i 3); done
```

---

This script changes the settings for the first 16 ports at the card in Slot 2 to mode 3, which is specified by the 3 in the script. If a 32 or 48-port card is used, you must replace 16 with 32 or 48. The slot also must be changed in both positions in the script. For a 32-port card in slot 3, it looks as shown in Example 8-4, and for a 48-port card in slot 4, it will look like Example 8-5.

*Example 8-4 Script to change the fill word for a 32-port card in slot 3*

---

```
for ((i=0;i<32;i++)); do (echo 3/$i;portcfgfillword 3/$i 3); done
```

---

*Example 8-5 Script to change the fill word for a 48-port card in slot 4*

---

```
for ((i=0;i<48;i++)); do (echo 4/$i;portcfgfillword 4/$i 3); done
```

---

8. Enter **logout**, and close the window.

The fill word is now changed, and the Director will use the configured mode.

### 8.3.9 Changing buffer credits

Buffer credits are used as a flow control mechanism to avoid overwriting stored frames received in a port. In most cases, extended distances require that the number of buffer credits be increased at the port. Refer to “Buffer credits” on page 51 for more information.

**Note:** Changing F-Port BB credits is not recommended unless advised by your vendor. Setting BB credits on E-Ports must be done using the port configuration wizard.

The default is 8 buffer credits per F-Port. Follow these steps to change the buffer credits for an F\_Port:

1. In the IBM Network Advisor server or client, right-click the Director, and select **Telnet**. A new window opens. (You can also use an ssh connection to the Director, if needed.)
2. Enter the user name (in this case, admin), and press Enter.
3. Enter the password, and press Enter (the default is password).
4. Enter, for example, **setcontext 1** if you need to connect to a logical switch with FID 1.
5. Disable the port by entering: **portdisable 1/9**.

This command will change the buffer credits for a given Port Index (PI) 9 to 80:

```
portcfgfportbuffers --enable 1/9 80
```

This command will change the buffers for Port Index (PI) 9 to default (Buffer Credits = 8):

```
portcfgfportbuffers --disable 1/9
```

6. Enable the port by entering: **portenable 1/9**.
7. To display the assigned buffers use the **portbuffershow 1** command. This command shows the assigned buffers for each port at the card in Slot 1 and the remaining buffers for the port group.

To change the buffer credits for an F-port using the Element Manager of the Director:

1. Go to the port configuration window of the Director's Element Manager. Right-click the Director in the IBM Network Advisor, and select **Element Manager** → **Ports**. The Port Administration window is displayed.
2. Click **Show Advanced Mode** in the upper-right corner to see the configuration options.
3. Select the port you want to configure, and click **Disable**.

**Important:** This is disruptive for the traffic on the selected port.

4. Read the information at the confirmation panel, and click **Yes**.
5. Click **Edit Configuration** for the selected port, as shown in Figure 8-44 on page 283.
6. Click **F-port BB Credit**, and enter the required value of Buffer Credits, as shown in Figure 8-46.
7. Click **Enable** to enable the port after the configuration is done.

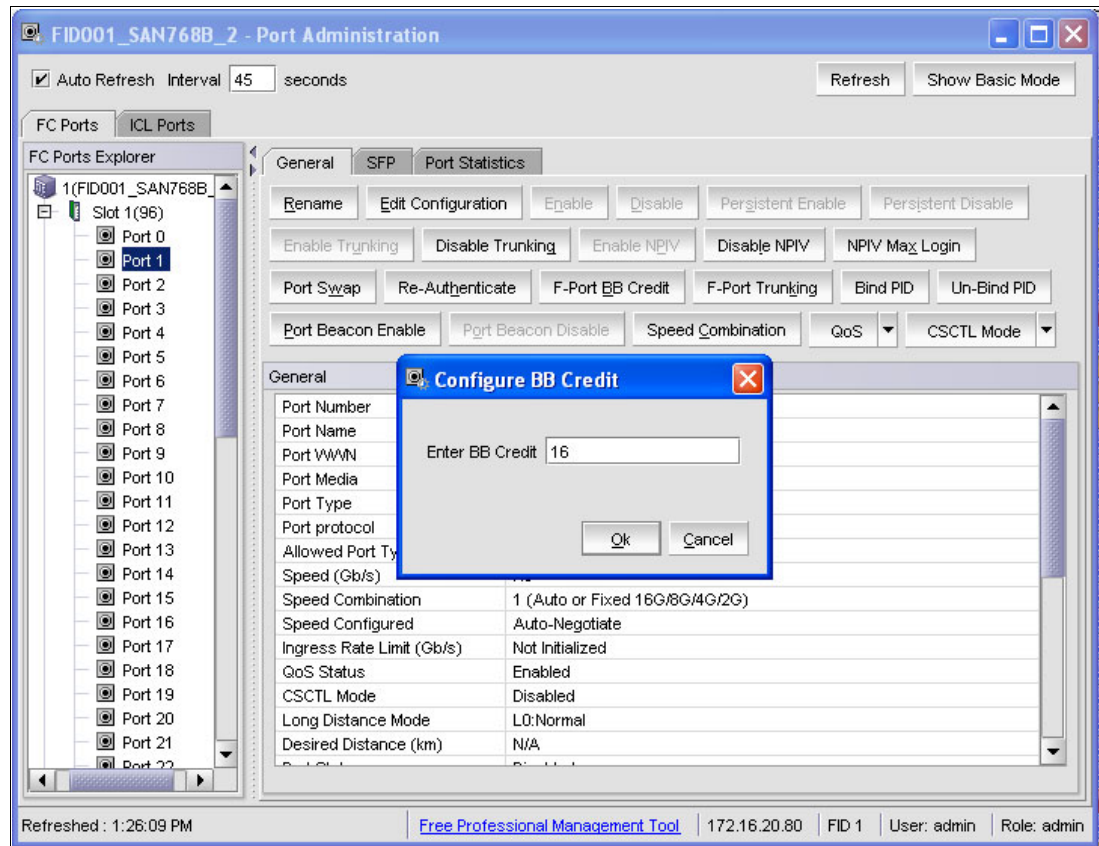


Figure 8-46 Buffer Credit

To change the buffer credits for an ISL, use the Element Manager of the Director as follows:

1. Follow the steps in “Changing the port type and speed” on page 282.
2. Open the drop-down menu for the Long Distance Mode, and select *one* of the following options:

**LE:<=10 km** (only use if the link is at maximum 10 km long). A total of 86 buffers will be allocated for a 16 Gbps link and 46 buffers will be allocated for a 8 Gbps link.

**LD:Auto** and enter the *maximum* link length in the Desired Distance (km) field. The buffer credits will be calculated during link initialization of the E\_Port and automatically allocated. The value you enter in the Desired Distance (km) field will determine the number of buffers allocated to the port based on a 2 KB frame size.

**LS:Static** and enter the link length in the Desired Distance (km) field. The buffer credits will be set for the given distance and the selected speed.

For an explanation of how buffer credits work, refer to “Buffer credits” on page 51.

3. Click **Next** after all selections are made.
4. At the Confirmation window, click **Save** after verifying the settings.
5. Click **Close** after you receive the configuration complete message.
6. Click **Enable** to enable the port after the configuration is done.

### 8.3.10 Setting up the Allow/Prohibit Matrix (optional)

To prohibit traffic between selected ports, you can use the Allow/Prohibit Matrix. To use the Allow/Prohibit Matrix, zoning must be active. If you use the default zone, create a new zone with all ports, as shown in 8.3.11, “Setting up zoning (optional)” on page 290.

To change the Allow/Prohibit Matrix:

1. Go to the Switch Admin window in the Element Manager of the Director. This can be done using the IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. Click **Show Advanced Mode** in the upper-right corner.
3. Select the **FICON CUP** tab.
4. Select the **CUP Port Connectivity** tab in the lower-left corner of the FICON CUP tab. The window shown in Figure 8-47 on page 289 opens.
5. Perform one of the following actions:
  - To make changes directly, select the Active Configuration and click **Edit**.
  - To create a new matrix, click **New**.
  - To copy an existing configuration, click **Copy**. You must provide a new Name and Description in the pop-up window, and then click **OK**.



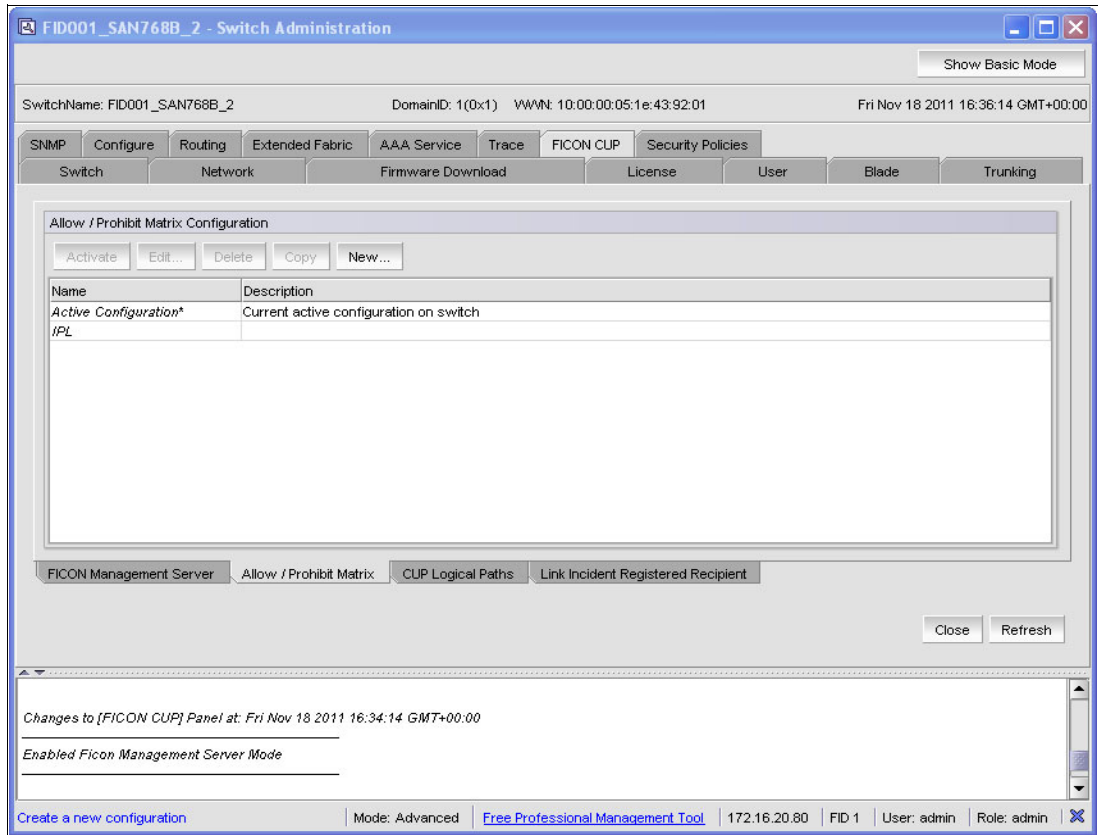


Figure 8-47 Allow/Prohibit configurations

6. After you make your selection and click Edit or New, the window shown in Figure 8-48 on page 290 appears.

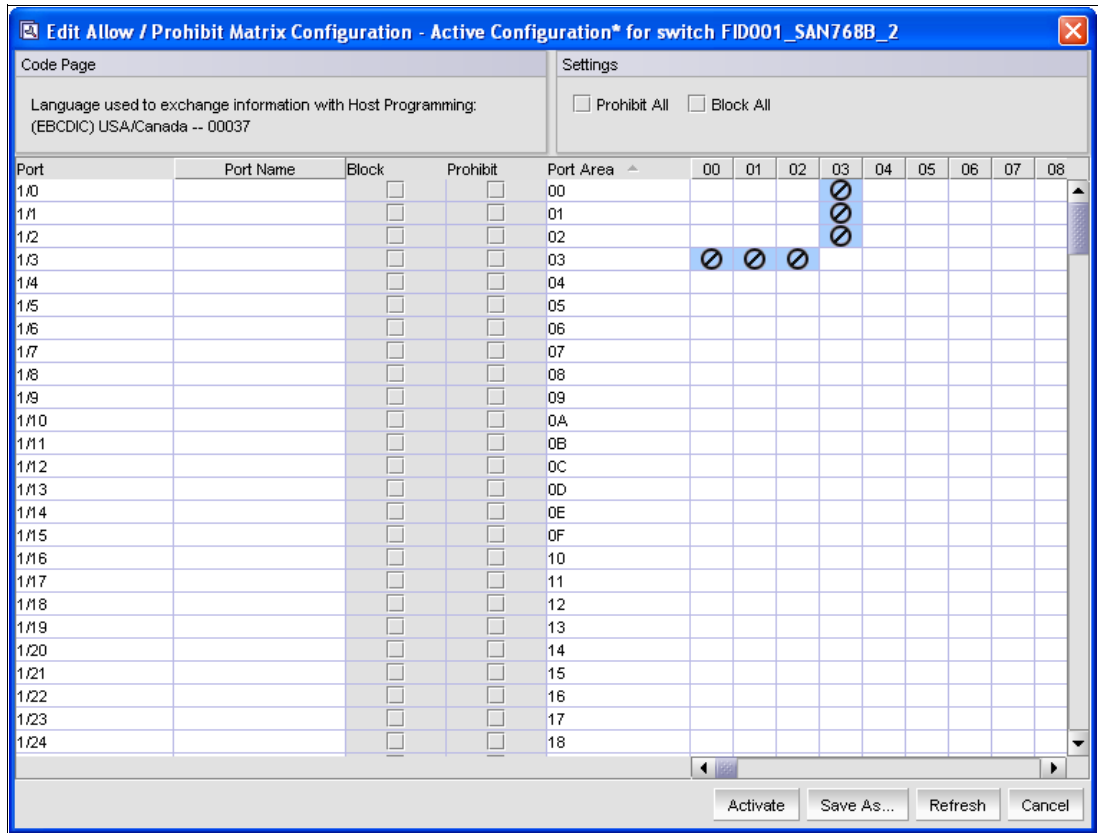


Figure 8-48 Allow/Prohibit Matrix

At this point, you can modify the matrix by clicking in the fields to allow or prohibit a connection. Figure 8-48 shows that port 03 (hex) is prohibited to connect to ports 00, 01, and 02. All other ports are allowed to connect to port 3.

Note the following points:

- The Block button will block the selected port. It is the same function as a port disable.
- The Prohibit button will prohibit all connections to and from the selected port.

If you do not want to activate the changes immediately, click **Save** (not available for the active matrix), or **Save As**, to create a new configuration.

7. After making the desired changes, you can activate the changes by clicking **Activate**.
8. Click **Yes** to confirm the action.

### 8.3.11 Setting up zoning (optional)

Typically all ports in a fabric are configured into a single large zone. In some cases, usually in cascaded environments or installations with different applications running, multiple zones are used to ensure that no unintended traffic or unnecessary RSCNs are transmitted by the fabric.

Although all access can be used for FICON environments, prohibits require at least one zone to be activated. Even if you do not plan to use the prohibit/allow matrix, configuring a single zone with all ports in it accomplishes everything “All Access” accomplishes while providing the flexibility to use prohibits or more restrictive zoning in the future.

When mixing operating systems in the same fabric, it is important to separate the traffic from different operating systems. The recommended best practice is to use virtual fabrics to keep this traffic separate. If using virtual fabrics is not an option, a more granular zoning plan must be implemented to isolate operating systems and the storage they use.

Although disk and tape mirroring uses FCP, it does not have to be isolated from FICON but the best practice is to isolate it, preferably in its own virtual fabric.

To set up zoning:

1. In the IBM Network Advisor, select **Configure** → **Zoning** → **Fabric** as shown in Figure 8-49.




Figure 8-49 Setting up Zoning

Figure 8-50 on page 292 is displayed, and in the Zoning Scope drop-down menu in the upper-left corner of the window, select the Director for which you want to create Zoning.

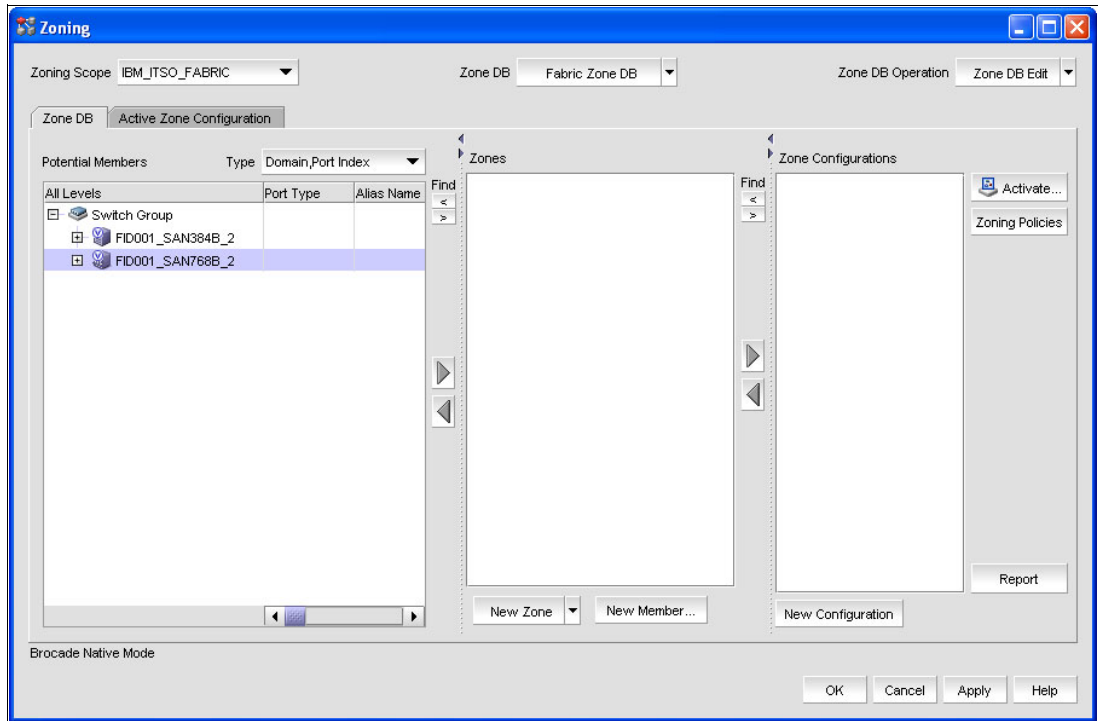


Figure 8-50 Zoning dialog overview

2. To create a new Zone, click the New Zone drop-down menu, and select **New Zone**. Enter a name for the new Zone (in this case, Zone\_Sample).
  3. Ensure that **Domain Port Index** is selected in the Type drop-down menu, because port-zoning is used in FICON environments.
  4. Select the ports you want to include in the zone you just created at the left side of the window. (You can select more than one port at a time by pressing and holding the Ctrl key on the keyboard while selecting the ports with the mouse. If you want to select a range of ports, use the Shift key.)
  5. To place the selected ports in the new zone in the middle of the window, click the arrow pointing to the right. Repeat steps Figure 8-50 to 5 for all zones, you want to create.
- The result is shown in Figure 8-51 on page 293.




Figure 8-51 Creating zones

6. After all zones are created, you must create or put them in a Zone Configuration. To create a new Zone Configuration click **New Configuration**. Enter a name for the new Zoneconfig (in this case, ZoneConfiguration\_Sample).
7. Select the zones in the middle of the window that you want to place in the newly created Zone Config, and then click the arrow pointing to the right. The result is shown in Figure 8-52 on page 294.

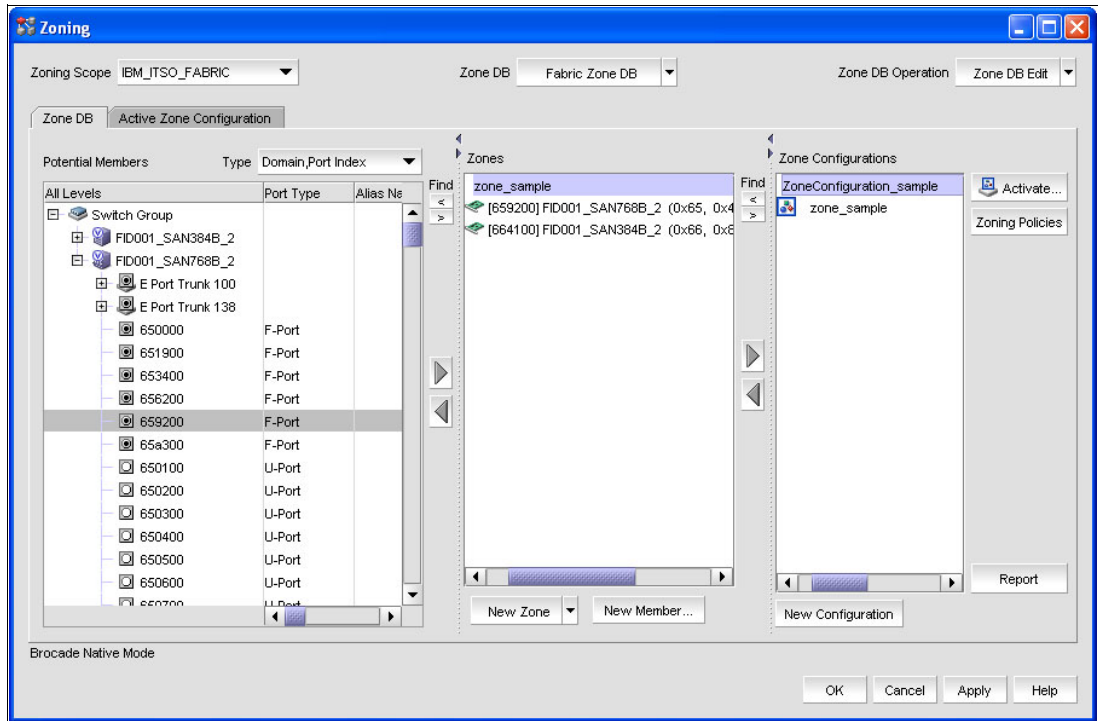


Figure 8-52 Creating a Zone Config

8. Now you must activate the new Zone Configuration. Select the Zone Configuration you want to activate (in this case, **ZoneConfiration\_Sample**), and click **Activate**.
9. You will reach the Activate Zone Config window. Verify that all ports are in the correct Zones. Click **OK** after you verify that all zones will activate the newly created Zone Configuration.
10. Click **Yes** when you reach the confirmation window.
11. After activation you will see that all active zones and the active Zone Configuration have green markers attached, as shown in Figure 8-53 on page 295.




Figure 8-53 Zoning activated

To verify the active zone, click the **Active Zone Config** tab. You will see the active Zone Config and the active zones with all zone members.

12. To leave the Zoning dialog, click **OK**.

13. If you get a confirmation window, read the warning, and click **OK**.

Zoning is now completed and active in the Director.

### 8.3.12 Configuring Port Fencing (optional)

Port Fencing is forced a port offline if a specified error rate is reached. A bad link will cause error recovery on the attached device, which can lead to degraded performance. If the link is an ISL, this will lead to recovery for many devices.

When Port Fencing is configured, the port is set offline automatically to keep the fabric stable. After the link is checked or repaired, the port must be manually set online.

To configure Port Fencing:

1. In the IBM Network Advisor, select **Configure** → **Fabric Watch** → **Port Fencing** from the menu bar, as shown in Figure 8-54 on page 296.

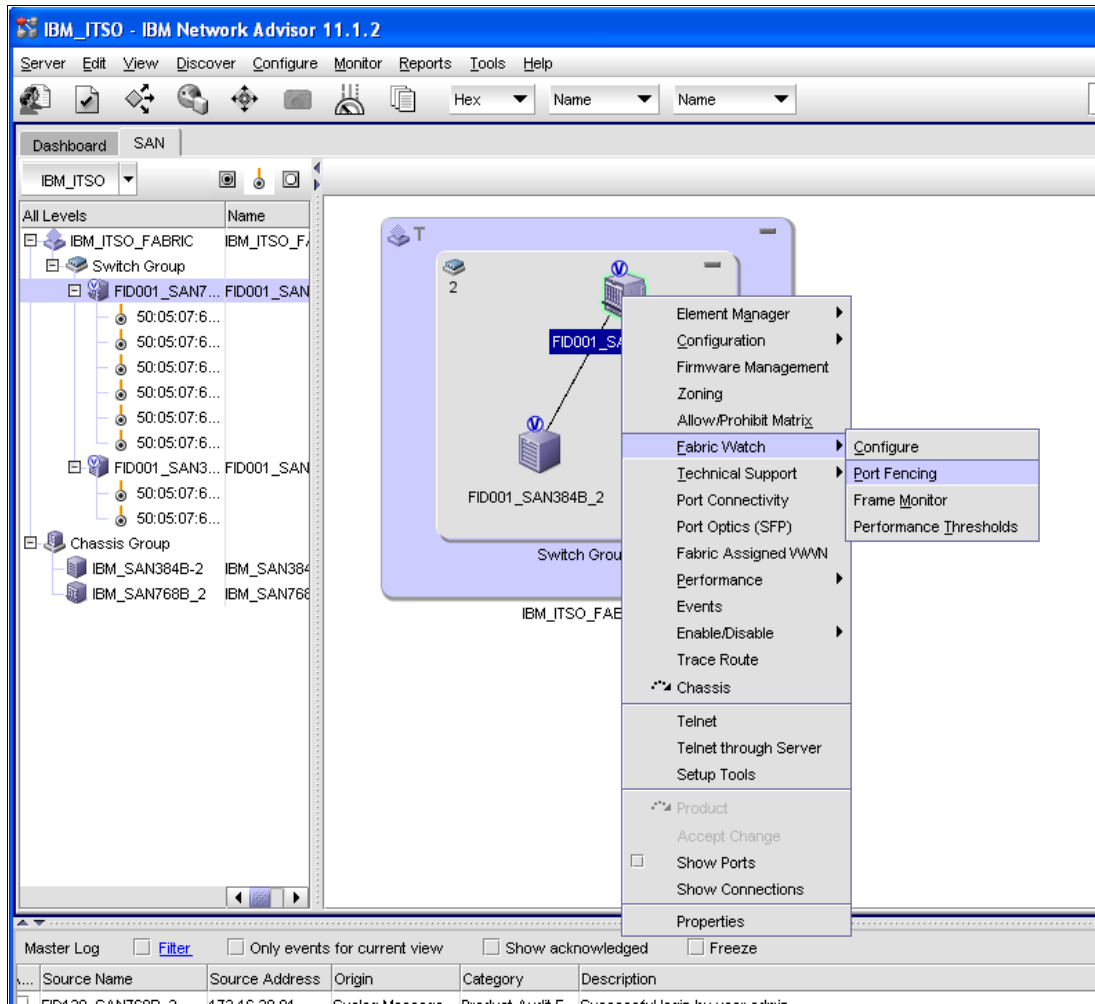


Figure 8-54 Port Fencing

2. Select the Fencing Policy from the Violation Type drop-down menu shown in Figure 8-55 on page 297 (in this case, the **Invalid CRCs** policy).



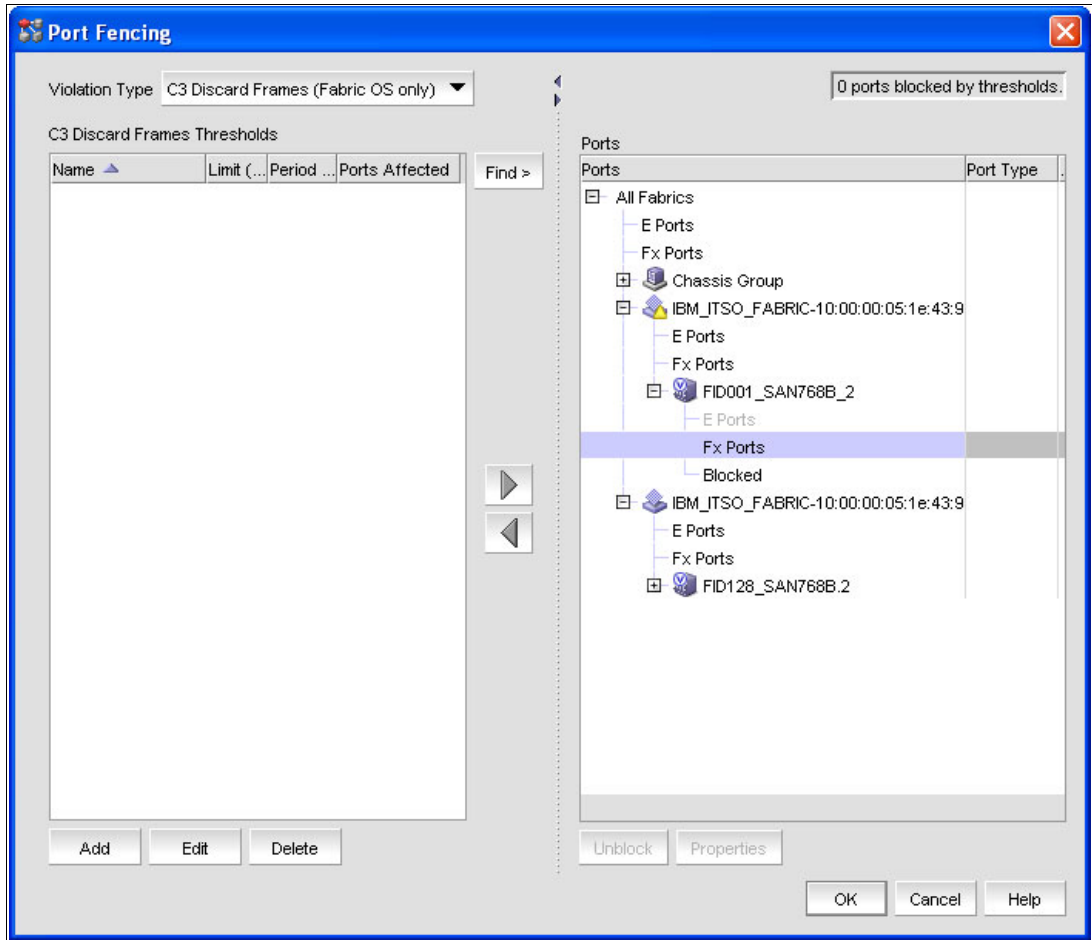


Figure 8-55 Port Fencing dialog

3. Select the policy on the left part of the window, or click **Add** to create a new one. If you want to change or verify the policy, click **Edit**.
4. Two selections are presented: m-EOS (which is for McData Directors) and FOS. In this case, we selected FOS for our SAN768B-2 Director.

There are two possible policy types in FOS: the Default and the Custom Policy. We selected the default policy, as shown in Figure 8-56.

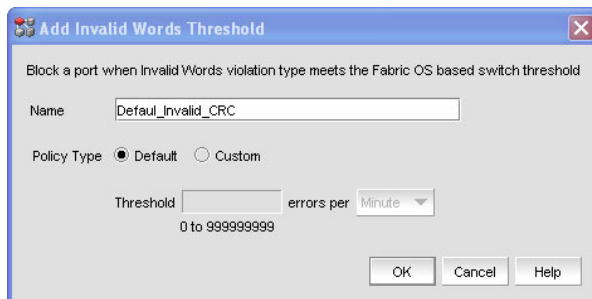


Figure 8-56 Default Policy

You can also change the policy name in this panel. (Some policies are only for FOS or m-EOS, as indicated at the selection from the drop-down menu.)

Click **OK** after the settings are done.

5. Assign the policy by highlighting it on the left site and highlighting the Director on the right side (or the fabric), and then clicking the arrow pointing to the right.

The policy is added to the Director as indicated by the green plus (+) sign as shown in Figure 8-57.

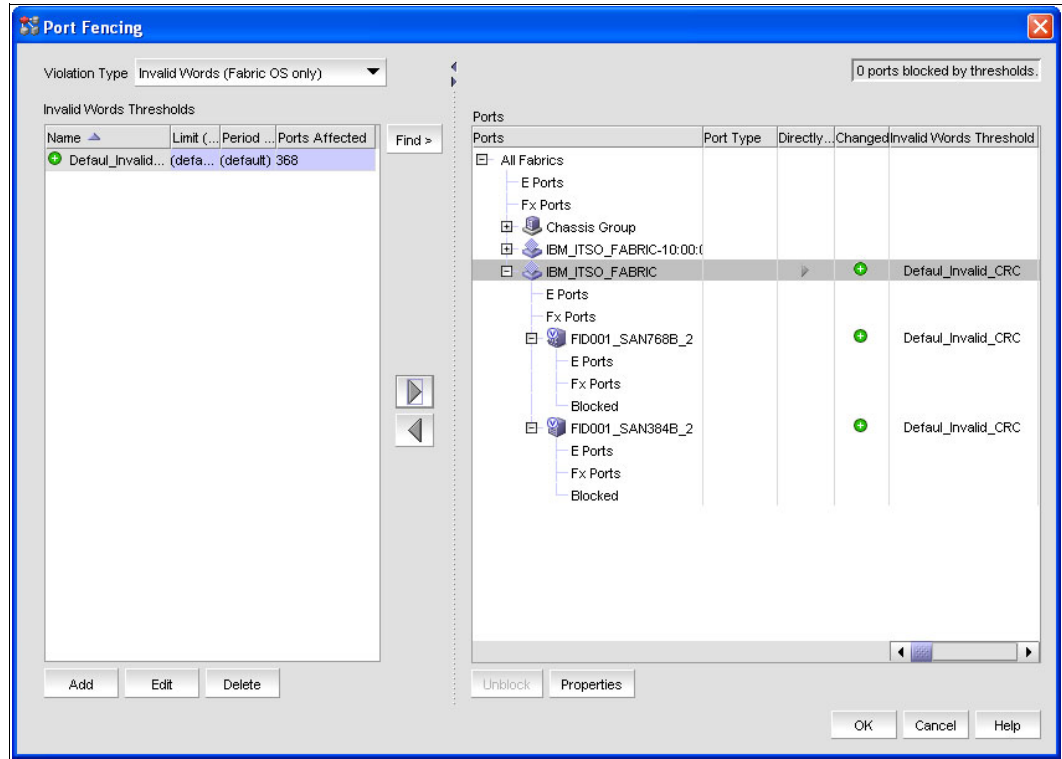


Figure 8-57 Port Fencing enabled

6. Click **OK** to activate the policy on the Director (or fabric).

You have now assigned the policy to the Director. When the threshold is reached for the assigned error type, the port will go offline.

Refer to *Fabric OS (FOS) Administration Guide* 53-1002148-03 and the *Fabric OS Command Reference Manual* 53-1002147-01 for other functions required for your installation, but not covered in this section.

## 8.4 Setting up a cascaded topology

A cascaded topology is a fabric configuration with two or more Directors but with no more than one hop (ISL) between a FICON channel (initiator) and a storage control unit (target). In a cascaded environment, the System z FICON channel must use the 2-byte addressing scheme, which includes the switch address, that corresponds to Port ID (PID) first byte (Director Domain ID) and the port address that corresponds to Port ID (PID) second byte (AREA) of the target device. The FICON Directors that will be merged have the same setup as any other FICON Director. Therefore, you must first complete the tasks in 8.3, “Setting up a FICON Director” on page 262 for both FICON Directors.

## 8.4.1 Setting up Inter-Switch Links

Inter-Switch Links (ISLs) connect the FICON Directors. Each ISL port in the FICON Director becomes an E\_Port type after initialization. Both sides of the ISL must have the same setup for them to function properly. This means that the port type (E-Port), port speed, Long Distance Mode, Trunking, and Quality of Service must be the same on both sides of the ISL. See 8.3.8, “Changing the port type and speed” on page 282 for more detailed information.

### Trunking (optional)

If trunking is enabled on both Directors, ISLs on the same port group automatically form a trunk. For trunking, all ISL link ports (maximum 8) must be in the same port group and must be set up with the same speed (for example, 16 Gbps). All fiber optic links in the trunk must be the same length (maximum of 30m apart).

If two of four ports have Quality of Service (QoS) enabled, two trunks are formed with two ports each (even if they are in the same port group): one trunk with QoS and one trunk without QoS. Trunking will only work in Brocade native Interoperability Mode.

To enable trunking:

1. Right-click the Director in the IBM Network Advisor, and select **Element Manager** → **Ports**.
2. In the Port Administration window, click **Show Advanced Mode** in the upper-right corner.
3. From the list, select the port that you want to enable for trunking.
4. Click **Enable Trunking**.

**Important:** This will directly enable trunking for the selected port. If this port is online, the link will go offline for a short moment, which will interrupt the traffic on the link.

### Traffic Isolation (TI) zoning (optional)

Using Traffic Isolation (TI) zoning, you can isolate traffic for some devices, thereby forcing them to use a specific ISL or trunk exclusively. If an ISL is selected that is a member of a trunk, the whole trunk is assigned to the TI zone. The trunk or ISL cannot be used by any other TI zone.

TI zoning is set up like regular zoning, but ISLs are included. To create a TI zone:

1. Perform steps 1 to 7 in “Setting up zoning (optional)” on page 290, but select **New TI Zone** instead of **New Zone**.
2. Add both sides of the ISL or trunk (E\_Port) and also F\_Ports to the new TI zone, which should be named with TI at the beginning of the zone name. The TI zone will not be included in the Zone Config on the right side of the Zoning window (see Figure 8-58 on page 300).

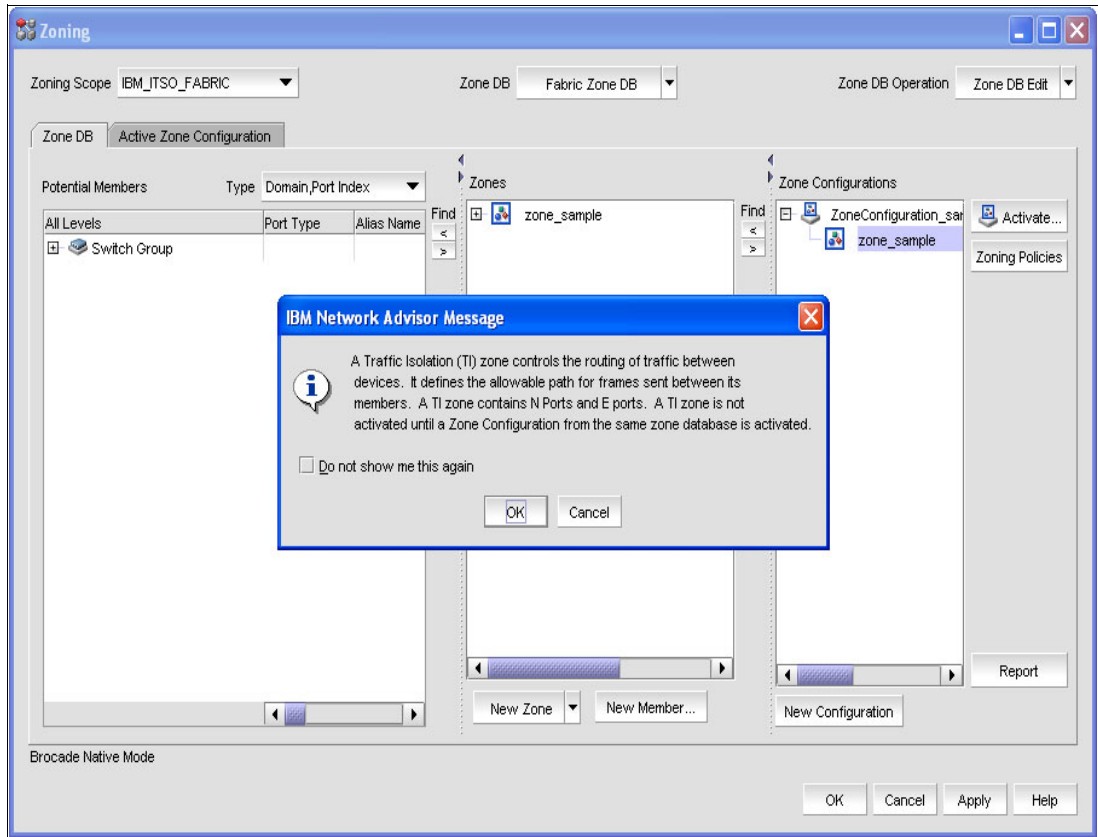


Figure 8-58 TI zoning setup

3. To activate the TI zone, right-click the TI zone, and select **Configured Enabled**. You can also select **Configured Failover**. If Failover is enabled, the devices in the TI zone will use another available ISL if the original ISLs in the TI zone fail (which is not in the TI Zone). When the configured ISL is available again, it will automatically fall back, as shown in Figure 8-59 on page 301.

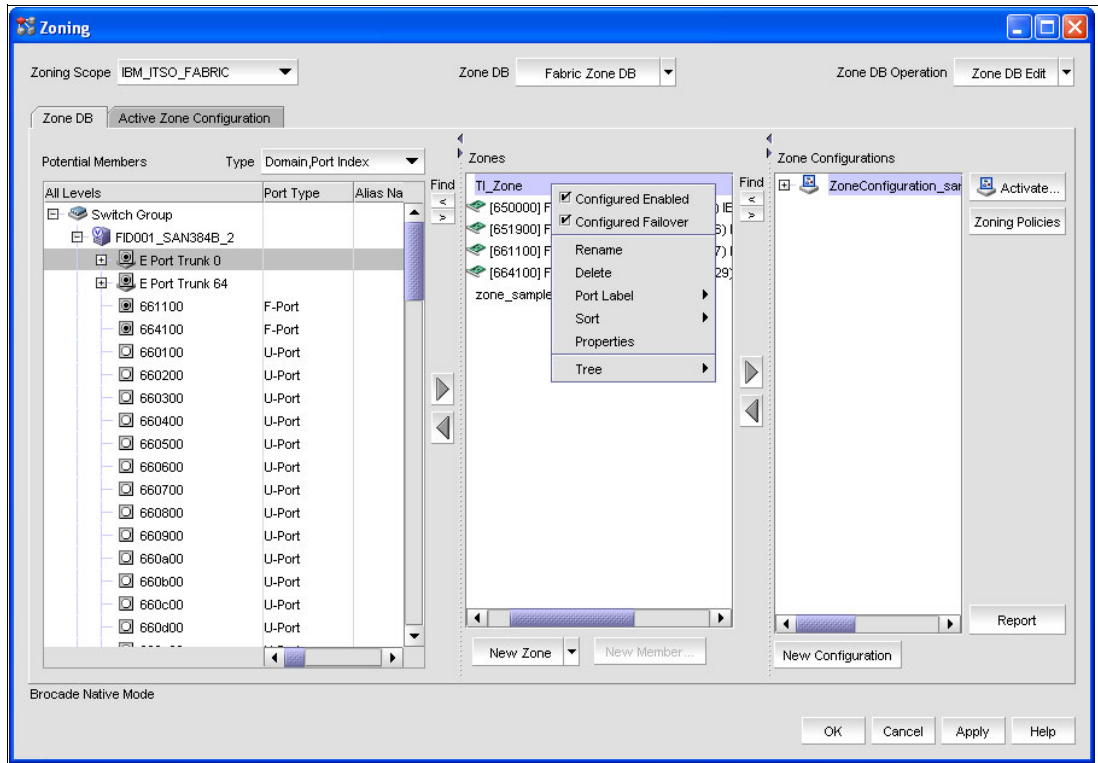


Figure 8-59 TI Zone activation

**Important:** When Failover is disabled, the devices in the TI zone will not use another available ISL outside its TI zone.

4. To activate one or more created TI zones in the Directors, click **Activate** in the Zoning window.
5. Compare the settings at the confirmation window, and select **OK** to activate the zoning changes.
6. If a second confirmation window appears, read the message, and select **Yes**.
7. Click **OK** after the changes are activated.
8. To verify that the TI zone is active, click the **Active Zone Config** tabs shown in Figure 8-60 on page 302.

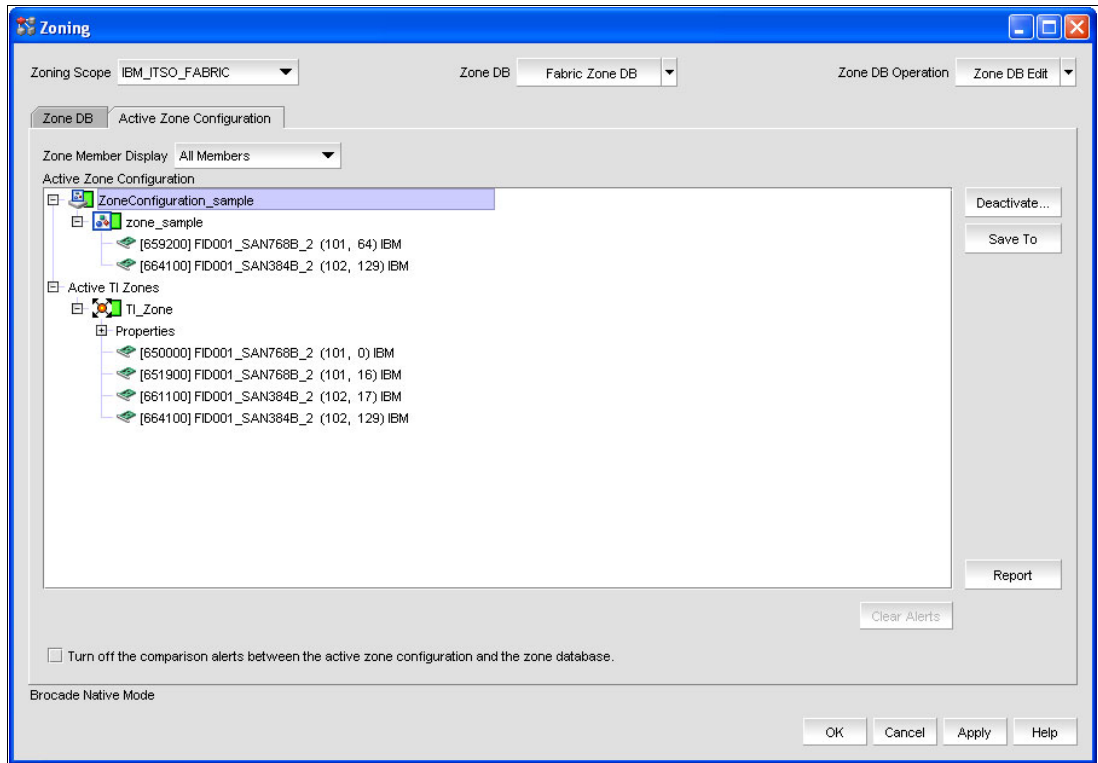


Figure 8-60 TI zoning activated

## Quality of Service (QoS) (optional)

Using Quality of Service (QoS), you can decrease the usable bandwidth for a device. This is done with the Ingress Rate Limit (Mbps) setup already shown in 8.3.8, “Changing the port type and speed” on page 282.

To enable QoS:

1. Right-click the Director in the IBM Network Advisor, and select **Element Manager** → **Ports**.
2. In the Port Administration window, click **Show Advanced Mode** in the upper-right corner.
3. From the list, select the port that you want to enable for QoS.
4. Click **QoS Enable** for the selected port.

There is also a possibility to use QoS zones. However, you must use WWN zoning instead of port zoning.

To use Quality of Service zones:

1. Complete the steps in “Setting up zoning (optional)” on page 290, but select **WWN** instead of Domain, Port Index in the Type drop-down menu.
2. Right-click the created zone, and select **QoS Priority** → **QoS\_High** or **QoS\_Low** as shown in Figure 8-61 on page 303. The default is **QoS\_Medium**, for all zones.

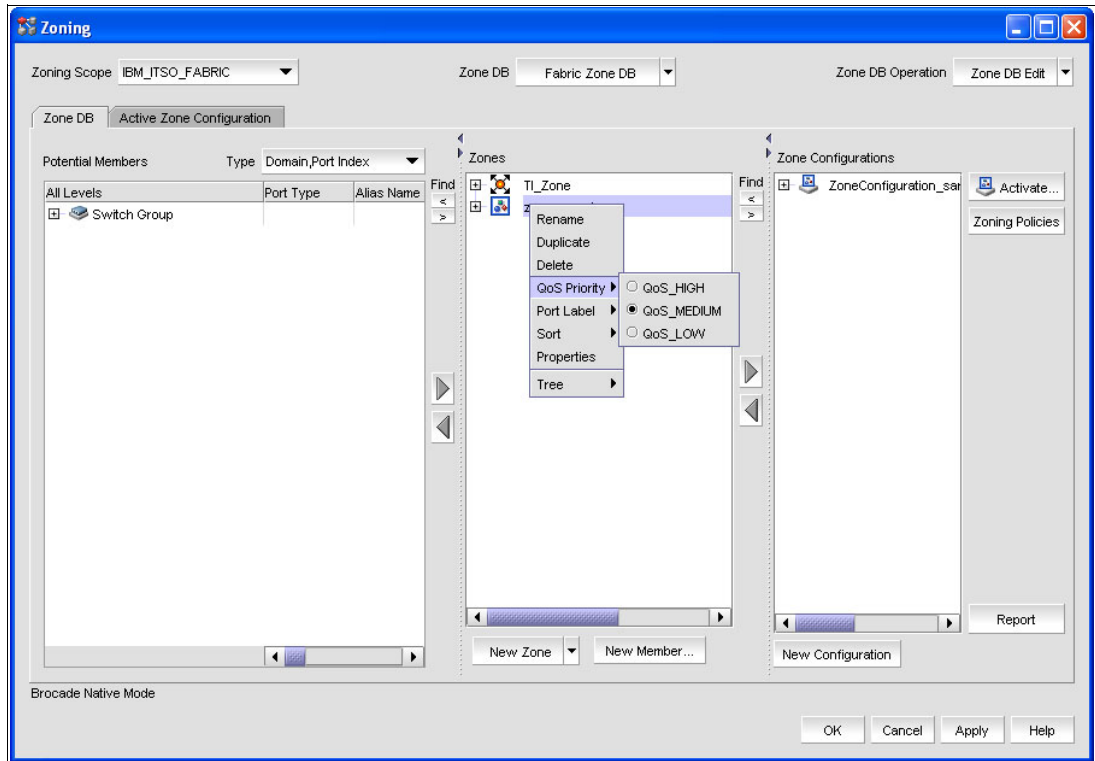


Figure 8-61 QoS priority selection for WWN Zoning

3. The name of the zone will be changed to QoSH (for high priority) or QoS\_L (for low priority) in front of the given zone name.
4. To activate the QoS zone, click **Activate** in the zoning window.
5. Compare the settings at the confirmation window, and select **OK** to activate the zoning changes.
6. If a second confirmation window appears, read the message, and select **Yes**.
7. Click **OK** after the changes are activated.
8. To verify that the zone is active, click the **Active Zone Config** tab.

## 8.4.2 Setting up a high integrity fabric

To merge the two Directors, the Switch Connection Control (SCC) Policy must be set up to form a high integrity fabric. Therefore, you must set the policy to Strict, which will distribute the policy only to FOS 5.2.0 or later Directors, resulting in a stable and secure fabric. In the SCC policy, we show how to include the World Wide Node Name (WWNN) of the two Directors in the fabric.

**Note:** The FICON configuration wizard automatically configures the security policies described here, which are required for 2-byte addressing, even if not cascaded. To add a switch after using the FICON Configuration wizard, use the FICON Merge Wizard.

To do this, follow these steps:

1. Go to the Switch Admin window in the Element Manager of the Director. This can be done using the IBM Network Advisor by right-clicking the Director and selecting **Configure** → **Element Manager** → **Admin**.
2. Click **Show Advanced Mode** in the upper-right corner.
3. Select the **Security Policies** tab, as shown in Figure 8-62.

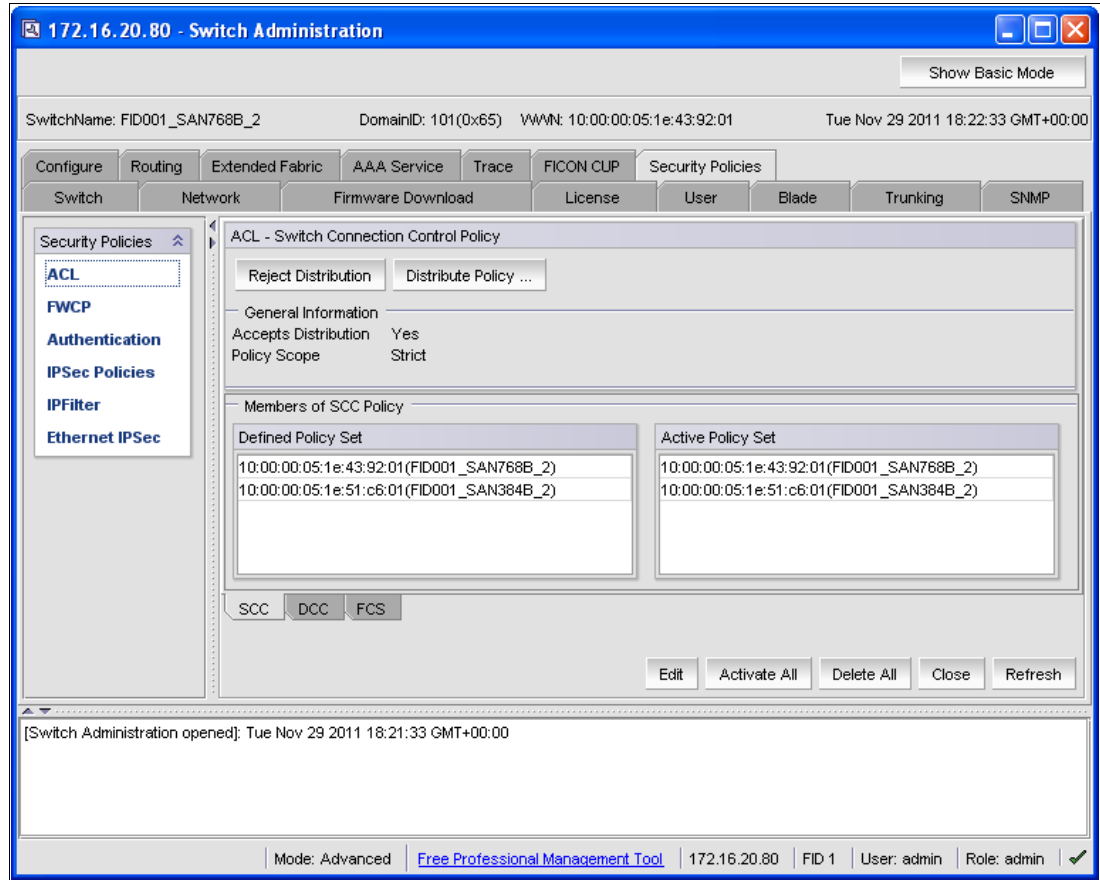


Figure 8-62 Security Policies

4. On the left side of the window, click **FWCP**, which is the Fabric Wide Consistency Policy.
5. Select **Strict** at the SCC Consistency Behavior drop-down menu, as shown in Figure 8-63 on page 305.



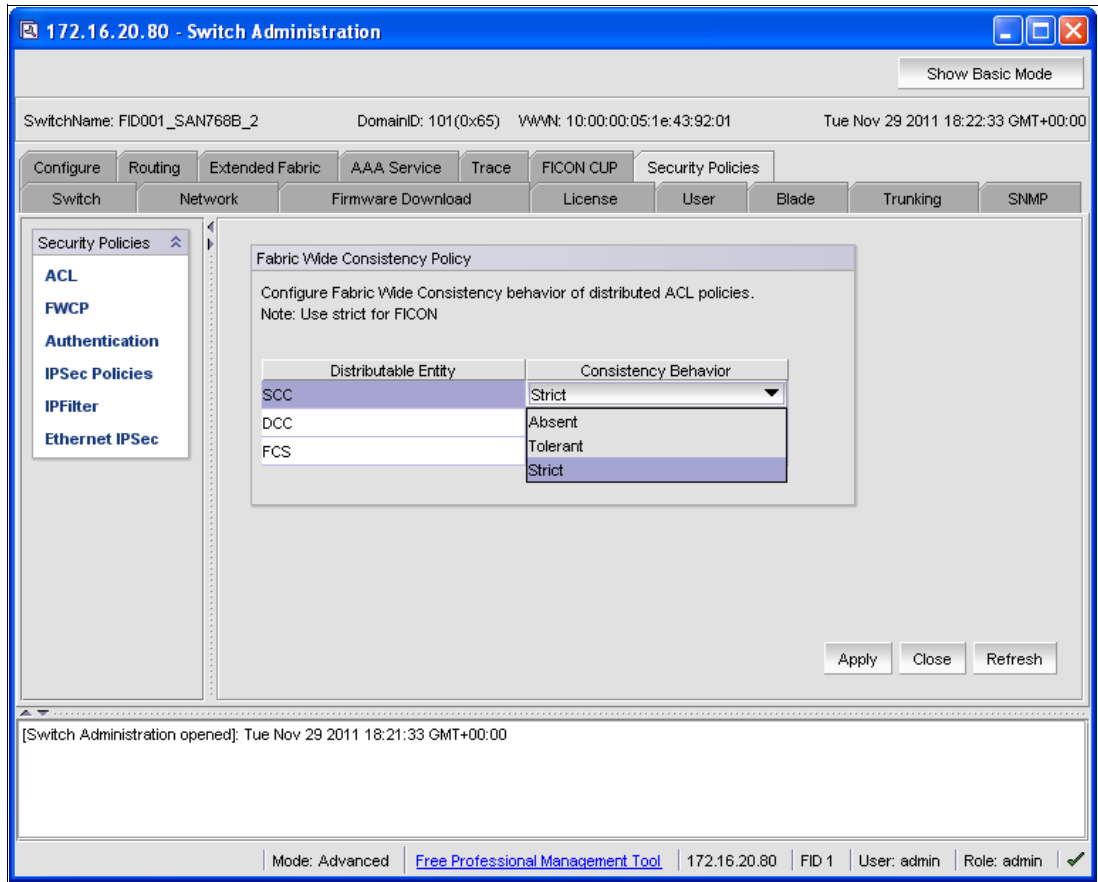


Figure 8-63 SCC Consistency Behavior

6. Click **Apply**.

**Note:** This will change the security settings of the Directors. From now on, each System z server must use the 2-byte addressing scheme to initialize the link.

7. Select **ACL** on the left side of the window.

8. Click **Edit** to set up the SCC Policy, as shown in Figure 8-64 on page 306.

9. Select only SCC, and click **Next**.

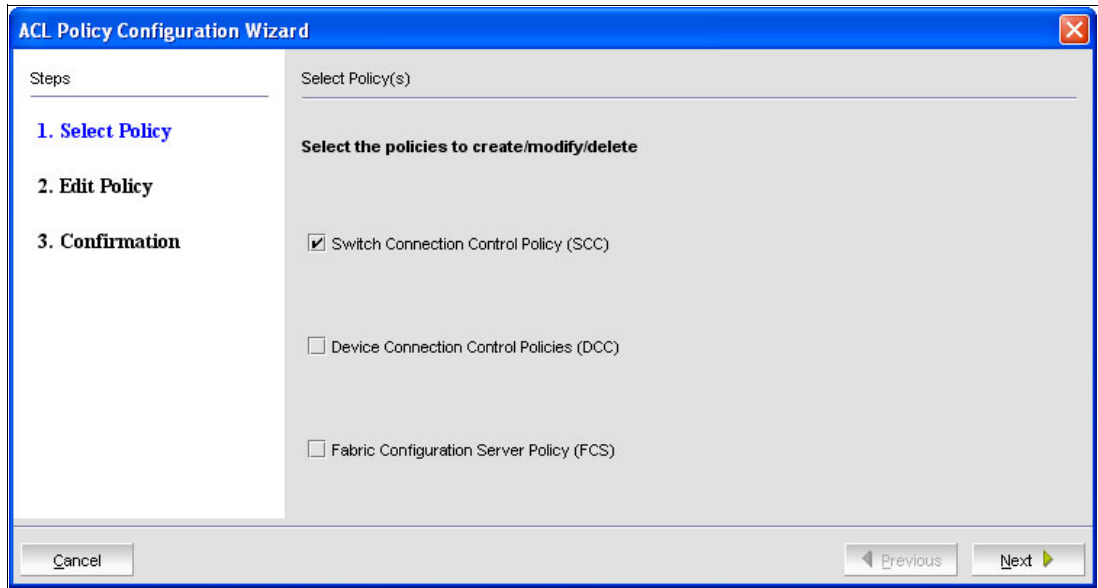


Figure 8-64 SCC edit policy

10. Click **Modify** to add the WWNNs of the two Directors shown in Figure 8-65.

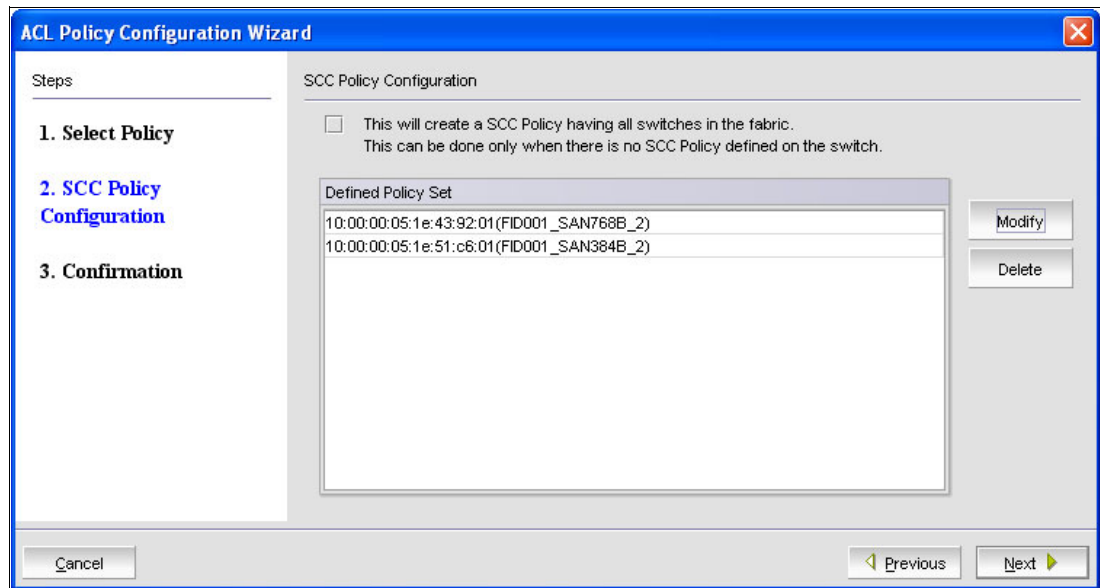


Figure 8-65 Modify WWNNs of Directors

11. Select the Director on the left side, and click **Add Switch** (see Figure 8-66 on page 307). This will add the WWNN of the Director you are connected to.

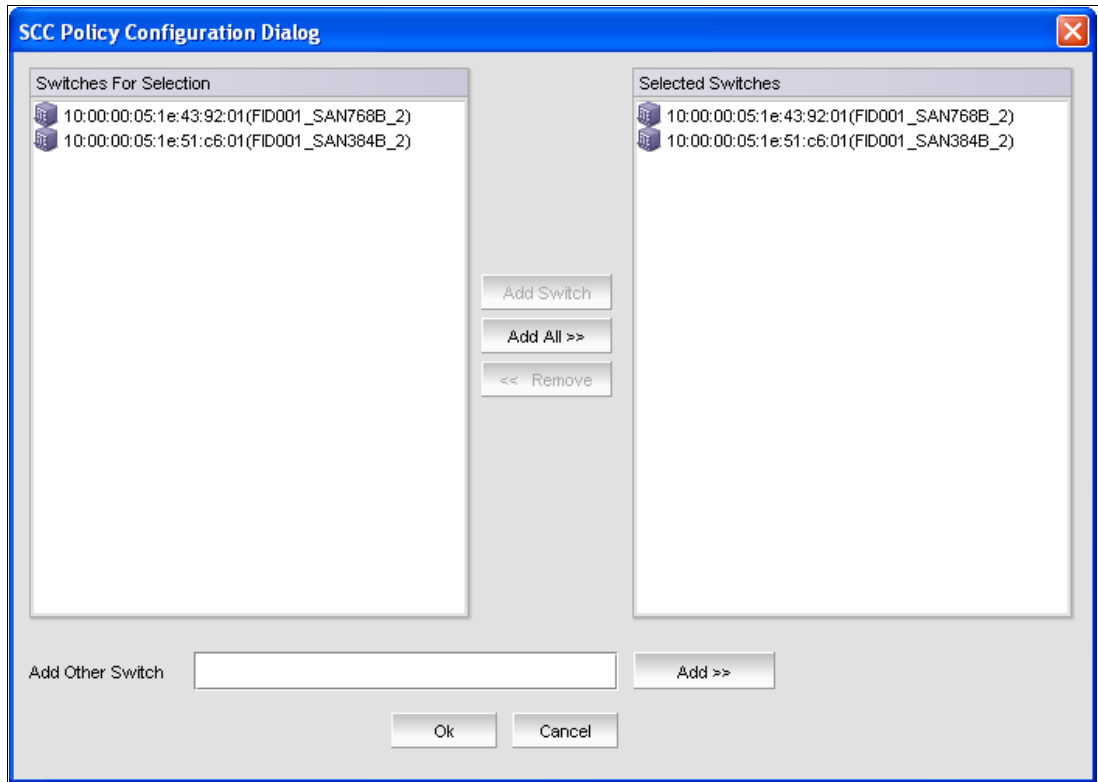


Figure 8-66 Creating SCC policy

12. In the Add Other Switch field, enter the WWNN of the second Director, and click **Add**. (The WWNN of the Director is shown on top of the Switch Admin window).
13. Click **OK**, and then click **Next**.
14. Verify the settings you just made, and click **Finish**.
15. Back in the Switch Administration window, you now see Defined Policy Set. Click **Activate All** to activate the Policy, as shown in Figure 8-67 on page 308.

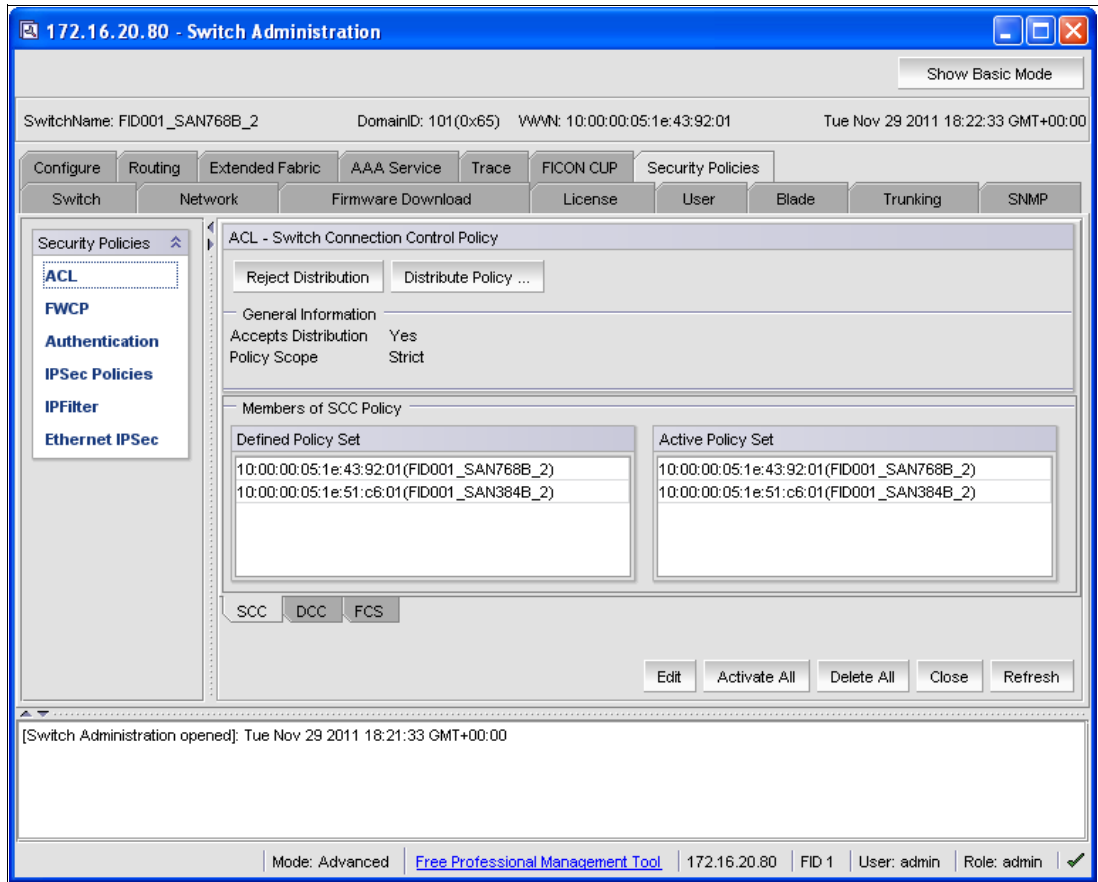


Figure 8-67 SCC policy activated

16. Perform steps 1 on page 304 through 6 on page 305 on the second FICON Director. (Set the Fabric Wide Consistency Policy for SCC to Strict.)
17. Connect the ISLs. The ISLs are initialized, and a cascaded fabric is created. Remember to set up the ISLs on both sides to the same settings as described in “Setting up Inter-Switch Links” on page 299.

The SCC Policy will be distributed to the second Director automatically.

**Note:** You do not need to edit the ACL list on the other Directors already merged on the fabric.

In IBM Network Advisor, you will now see that the two Directors are merged and connected through ISLs, as shown in Figure 8-68 on page 309.

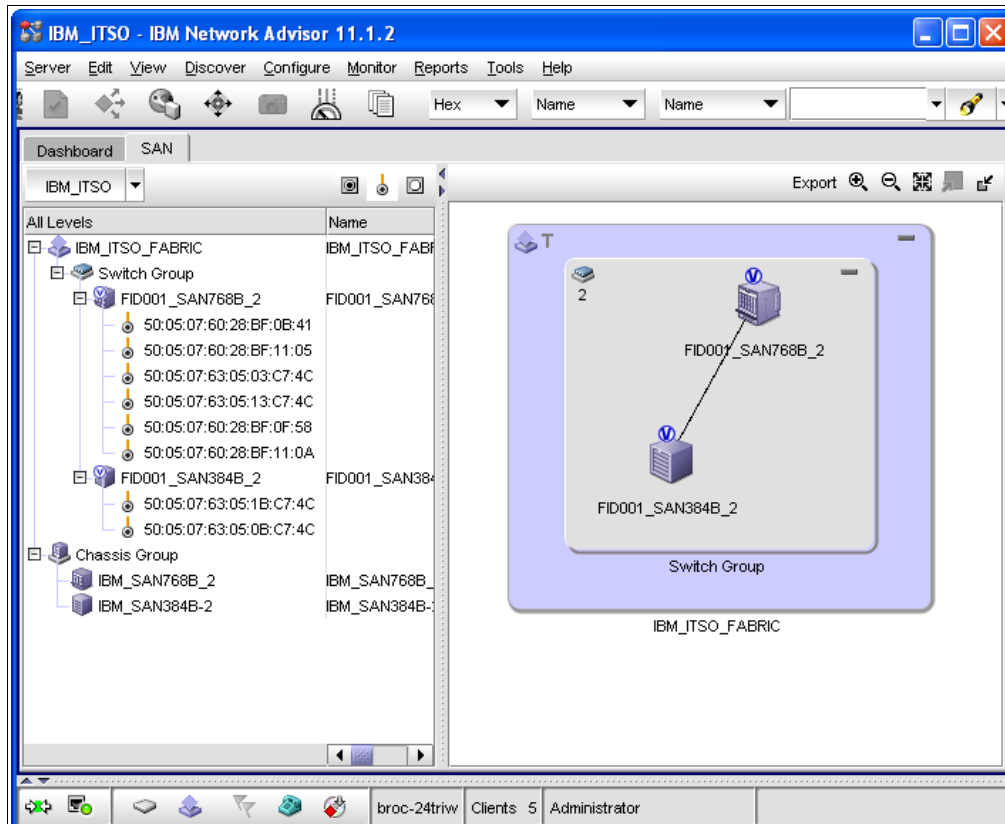


Figure 8-68 FICON Cascaded Director IBM Network Advisor overview

To verify the ISLs, right-click them in the IBM Network Advisor Directors View, and select **Properties**. You will see all active ISLs and their settings on both Directors, as shown in Figure 8-69 on page 310.

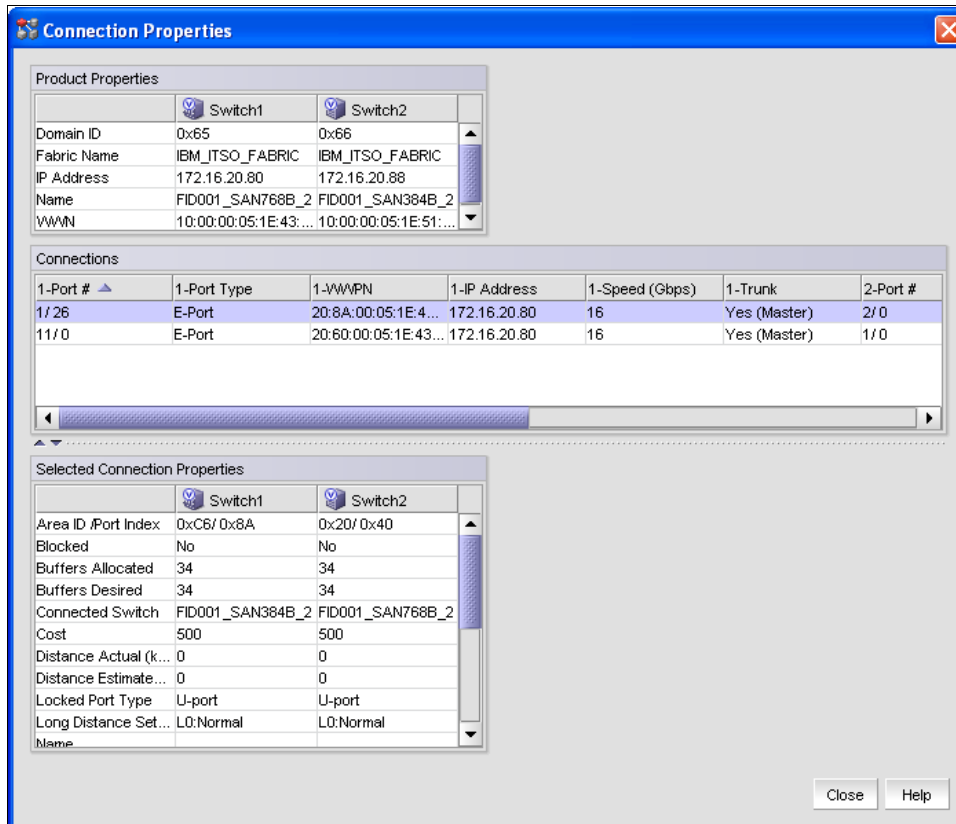


Figure 8-69 ISL properties window

If ISLs are missing, go to the Port Admin window of the two Directors and check the settings on both sides. Use the Advanced View to see all settings. Scroll to the right side, and check the Additional Port Info field (this field also displays some error reasons for the ISL).

## FICON Merge Wizard

After enabling the proper security requirements using the FICON configuration wizard or the procedure documented in the previous section, fabrics can be merged using the FICON merge wizard. Figure 8-70 on page 311 and Figure 8-71 on page 311 are examples of the merge wizard windows.

## Wizard

### Select Fabrics

Right click on one of the switches and select

FICON
Configure Fabric  
Merge Fabrics

---

Select Fabrics

Select two fabrics to merge for cascaded FICON setup.

| Available Fabrics                    | Selected Fabrics                     |
|--------------------------------------|--------------------------------------|
| IP62_7800_BJ_Fabric (172.16.20.62)   | IP62_7800_BJ_Fabric (172.16.20.62)   |
| PLUTO_1_DID_66_Fabric (172.16.20.88) | PLUTO_1_DID_66_Fabric (172.16.20.88) |

1. Select fabrics in "Available Fabrics" column.
2. Use arrow key to move to the "Selected Fabrics" column.
3. Click the "Next" button at the bottom right of the screen..

Figure 8-70 Select Fabrics

## Wizard

### Setup Merge Options

Always check FMS Mode, even if not using CUP

Set up merge options

Turn on FICON Management Server (FMS) mode on all the switches of the selected two fabrics. This feature must be enabled on a switch to allow Control Unit Port (CUP) management. If enabled, a Configuration is required for the PDCM to take effect.

FMS Mode  Enable      256 Area Assignment: Disable

Select one of the two fabrics of which AD, ACL, and zone DB, if defined, will be cleared.

sw20d0 (9.56.197.6)  
 sw20d1 (9.56.197.9)

Select the fabric that will merged in. All connections to the selected switch (in this example sw20d1) will be dropped..

Figure 8-71 Setup Merge Options

## 8.5 FICON Directors in an extended distance environment

In an extended distance environment, using a Wave Division Multiplexing (WDM) platform will require you to set up the Director to match the WDM configuration. Check which WDM platforms are IBM System z qualified and ensure that you are using the correct SFPs; otherwise, you can overdrive them in the WDM.

Configure the ISL ports as E-Ports (not auto-detect U\_Ports) and use a fixed speed (not auto-negotiate). Also, the Long Distance Mode must be configured for the correct distance. See 8.3.8, “Changing the port type and speed” on page 282 and 8.3.9, “Changing buffer credits” on page 286 for configuration information.

If Figure 8-72 appears when setting the port long distance parameters, do not consider it. This message is a mistake and will be corrected on an IBM Network Advisor future version.

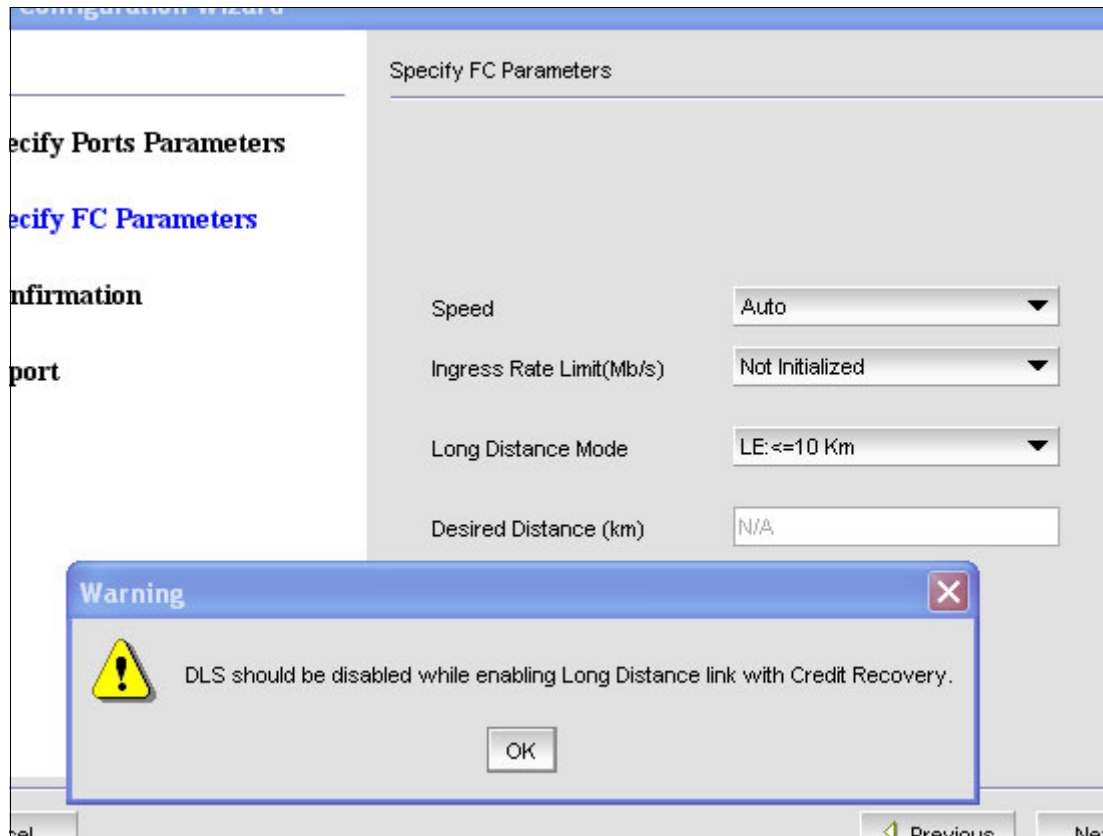


Figure 8-72 IBM Network Advisor pop-up error message

By default, the Director uses VC\_Link\_Init with ARB Fillword to initialize the ISL. However, some Dense WDM platforms do not support this mode. For those WDM platforms, you must use Idle Fillwords, or disable it and use ISL R\_RDY Mode. Also, you must not use Quality of Service and Credit Recovery in this mode. They must be disabled before Idle Fillwords are used, or ISL R\_RDY Mode is enabled. ISL R\_RDY does not support trunking.

The following commands are needed:

- ▶ **portcfgport slot/port**
- ▶ **portcfgspeed slot/port speed-level**
- ▶ **portcfgqos --disable slot/port** to disable QoS.  
**portcfgcreditrecovery --disable slot/port** to disable credit recovery.
- ▶ **portcfglongdistance slot/port LS 1 100** to enable Long Distance in LS Mode for 100 km with VC\_Link\_Init enabled for ARB Fillword (1).  
**portcfglongdistance slot/port LS 0 100** to enable Long Distance in LS Mode for 100 km with VC\_Link\_Init enabled for IDLE Fillword (0).
- ▶ **portcfgshow slot/port** to verify the settings.



Example 8-6 illustrates how to set up Long Distance Mode to 100 km (LS Mode = static) with ISL R\_RDY Mode and Idle as Fillword, which is required for some WDM platforms.

*Example 8-6 DWDM example for Port 1/15*

---

```
portcfgqos --disable 1/15
portcfgcreditrecovery --disable 1/15
portcfglongdistance 1/15 LS 0 100
portcfgshow 1/15
```

---

## 8.6 FICON Directors in an intermixed environment

If you use FICON and Fibre Channel in the same Director, we strongly recommend that you create two logical switches: one logical switch for FICON and the other for Fibre Channel. For configuration details, refer to 8.3.4, “Setting up a logical switch (optional)” on page 270.

This action allows you to configure a FICON logical switch with all required settings and a Fibre Channel logical switch with its own settings.

## 8.7 Channel swap

A channel swap is useful if you need to move a connected device to another physical port in the Director without changing the definition in the System z server or the zoning. The port will have the same address as it had before the swap.

To perform a channel swap:

1. Right-click the director in the IBM Network Advisor, and select **Element Manager** → **Ports**.
2. In the Port Administration window, click **Show Advanced Mode** in the upper-right corner to see all possible functions.
3. Disable both the port you want to swap and the port you want to swap to.

**Note:** Disabling a port is disruptive.

**Note:** If you have intermittent errors to a System z server, set the attached CHPID offline before disabling the port on the Director side. Set the CHPID online after the swap is complete and the port is enabled again on the Director side.

4. Select the port you want to swap to and click **Port Swap**. Enter the Slot Number and the Port Number of the “defective” port and click **Swap**.
5. After the swap is complete, plug the cable to the new port in the Director.
6. Enable the port after the cable is connected and check the Node Descriptor for a valid connection, as shown in 10.8, “Node descriptor” on page 362. Set the port online in z/OS.

## 8.8 Backing up Director configuration data

We recommend that you have a backup in case of a disaster (the configuration file can be restored in a replacement Director). You should also perform a backup before a firmware upgrade. Before you back up Director, save the Director configuration data.

To save the Director configuration data:

1. Right-click any Director and select **Configuration** → **Save**.
2. In the Save Switch Configurations dialog, shown in Figure 8-73, select the Directors you want to back up, and add them to the Selected Switches list by clicking the arrow pointing to the right.
3. Click **OK** to save the Director configurations.

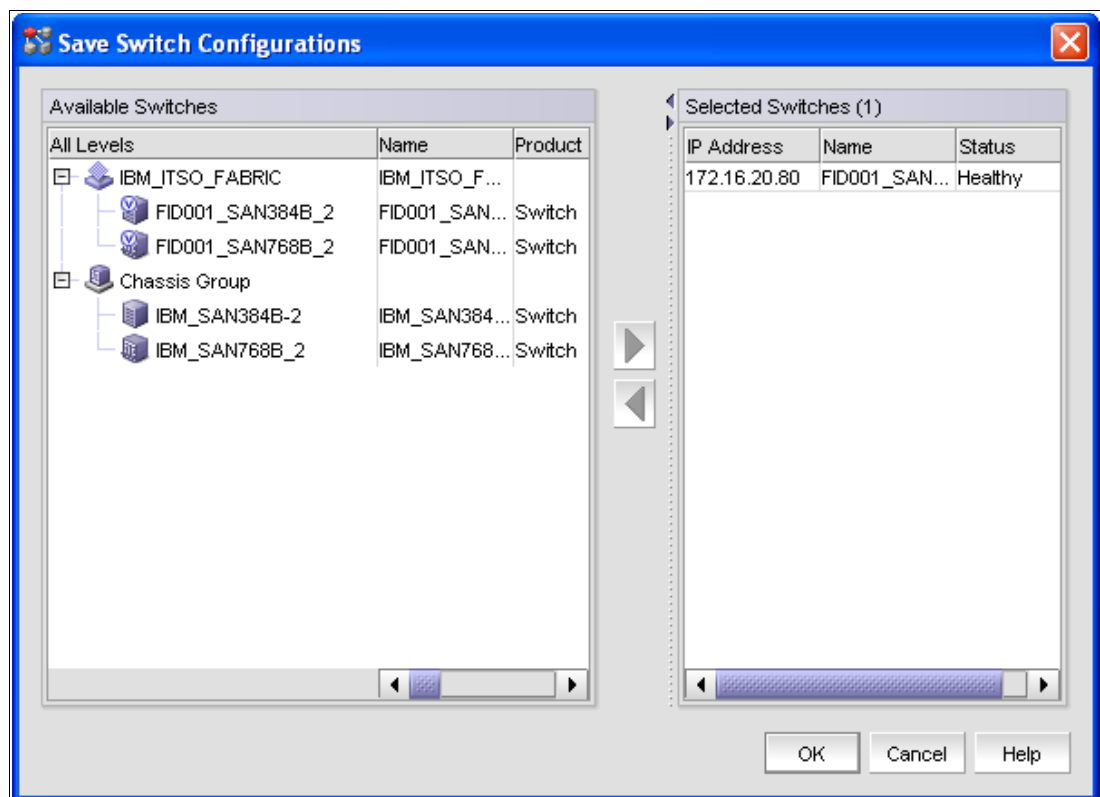


Figure 8-73 Save Switch Configurations

You can also schedule a backup by right-clicking any Director and selecting **Configuration** → **Schedule Backup**. Make the desired settings in the Schedule Backup dialog, and click **OK**. The backup will be performed automatically according to your settings.

You can view the backups by right-clicking any Director and selecting **Configuration** → **Configuration Repository**.

To restore, right-click the Director you want to restore, and select **Configuration** → **Restore**. You will obtain a list of available backup files. Select the backup you want to restore, and click **OK**.

**Important:** The restore of the Director's configuration data is disruptive because the Director will perform a reboot to activate the new configuration.

## 8.9 Backing up IBM Network Advisor configuration data

You must also backup the IBM Network Advisor configuration. Select **SERVER** → **Options** from the IBM Network Advisor menu bar. In the Options window, select **Backup** on the left and provide the necessary information. Click **OK** after the settings are done.

**Note:** We recommend that you save the backup on a network drive to be able to restore the data in case of a disaster or a hard disk failure.

The data can be restored through the IBM Network Advisor Server console (located in the program directory):

1. You must stop the IBM Network Advisor services by clicking **Stop**.
2. Select the **Restore** tab, and enter the location of the backup, or browse the filesystem. After you select the backup location, click **Restore**.





## Part 4

# Managing the FICON environment

This part discusses FICON operations and problem determination tools and techniques. Common issues are examined and descriptions are provided that explain how to resolve these issues using proven methods when working with FICON environments.





# Monitoring the FICON environment

In this chapter, we discuss verification and monitoring tools and techniques for FICON environments, which includes:

- ▶ HMC System Activity Display
- ▶ RMF Reporting
- ▶ IBM Network Advisor performance monitor
- ▶ Web Tools performance monitor

## 9.1 System Activity Display

The System Activity Display (SAD) on the HMC shows *individual* system activity for System z servers (also known as Central Processor Complexes or CPCs) or for a *group* of Central Processor Complexes. System activity includes the channel activity and physical processing activity that are defined in the system activity profiles, which are stored in the selected CPCs. For more information about assigning and customizing activity profiles, refer to *HMC Operations Guide*, SC28-6905.

Using the SAD, you can create a list of channels that you want to look at and monitor, if a performance problem is suspected to be caused by one or more FICON channels. Figure 9-1 shows a System Activity Display for CHPID usage on processor SCZP201. However, for performance troubleshooting, we recommend that you use the Resource Measurement Facility (RMF).

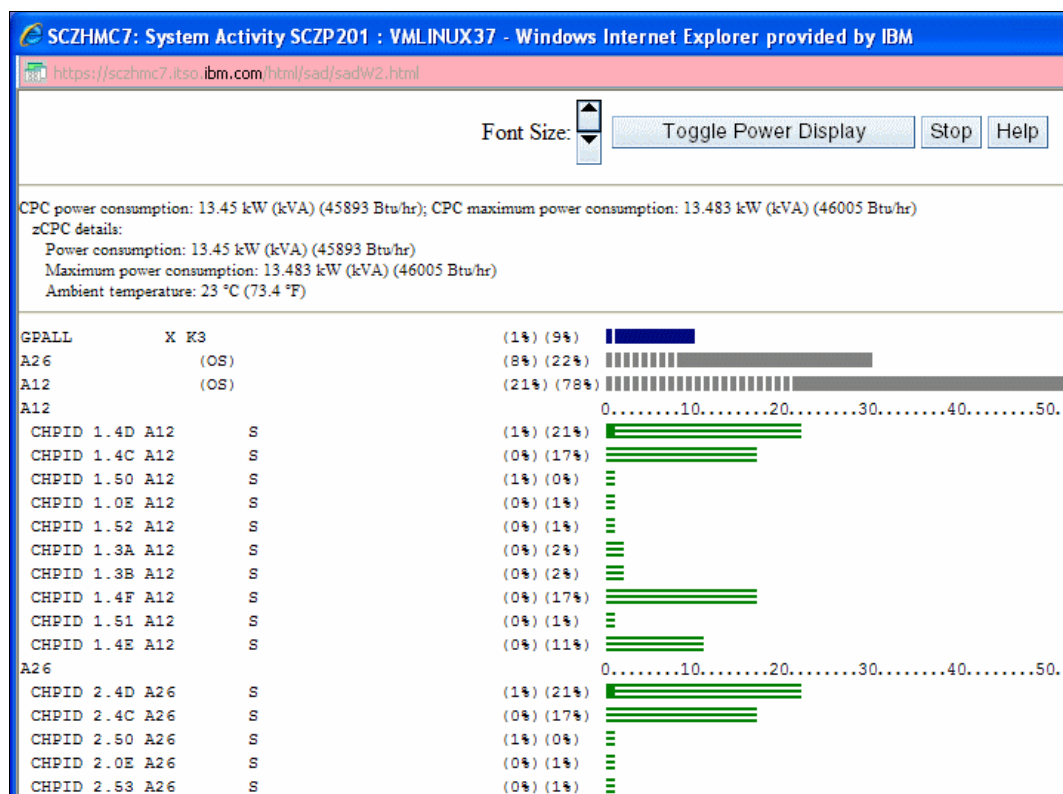


Figure 9-1 SAD frame displaying channel activity usage

**Note:** The utilization reported by the Activity task for most channel types will correspond with the utilization reported by the Resource Measurement Facility (RMF). For Fibre Channels, however, this task considers the channel to be busy any time an operation is pending, even if the channel is waiting for a device to respond. RMF, in contrast, looks at the amount of work performed versus the amount of work that *can* be performed by the channel. This means that if you have devices that are relatively slow to respond, leaving the channel waiting for a response but otherwise idle, activity will show a utilization that is significantly higher than that reported by RMF.



## 9.2 Resource Measurement Facility monitoring

Resource Measurement Facility (RMF) is a performance management tool that measures selected areas of system activity, including several varieties of the FICON channel, and presents the data collected in the form of System Management Facility (SMF) records. This data is essential for any kind of FICON channel performance monitoring.

For more information, refer to the IBM RMF web site at:

<http://www.ibm.com/servers/eserver/zseries/zos/rmf/>

## 9.3 Introduction to performance monitoring

Before you undertake any performance analysis, it is imperative to establish your performance objectives. The goal of performance management is to make the best use of your current resources to meet your objectives without excessive tuning efforts.

An essential ingredient in performance monitoring is understanding the components of response time. The measurements we calculate to report I/O response time and that RMF reports on are:

- ▶ Queue time
- ▶ Pend time
- ▶ Connect time
- ▶ Disconnect time

Figure 9-2 indicates the I/O performance measurement points. It depicts a configuration with multiple servers exploiting a cascaded FICON director environment; however, the measurement points are also valid for the other supported topologies (by excluding the PEND time for the second switch or both switches).

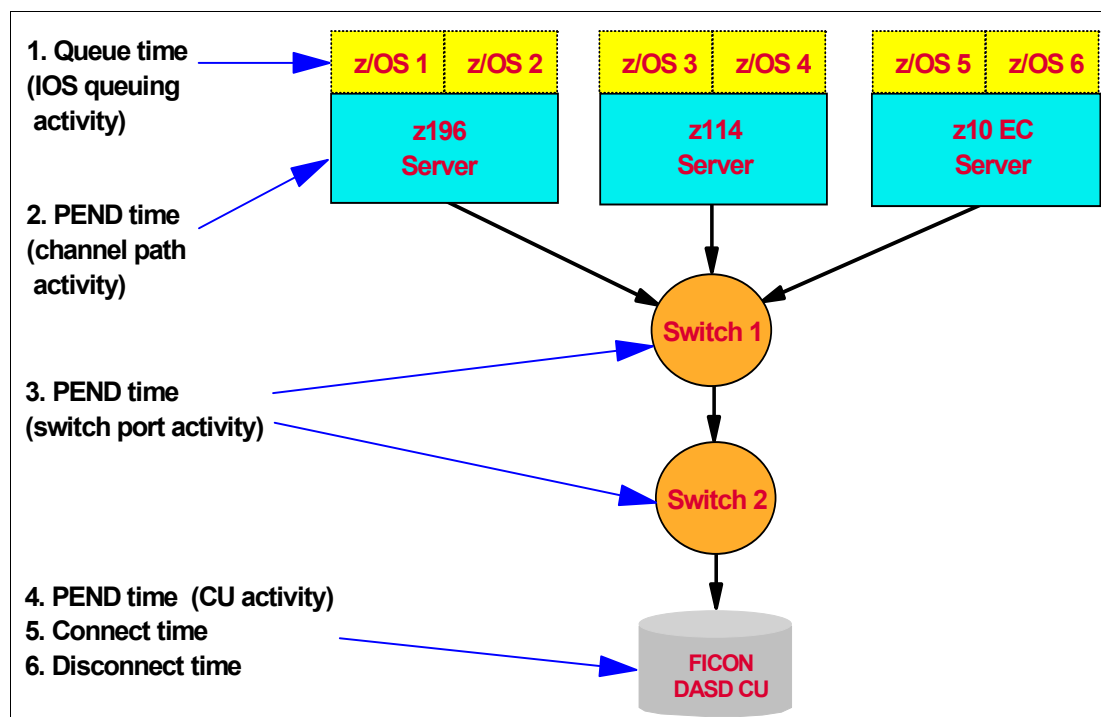


Figure 9-2 I/O performance indicators

FICON exploits CCW chaining, so there will be fewer channel ends and device ends. Also, PEND time ends at the FICON channel when the first command response (CMR) is received from the CU.

## 9.4 Introduction to Resource Measurement Facility

The Resource Measurement Facility (RMF) provides an interface to your System z environment that facilitates reporting and detailed measurements of your critical resources. RMF consists of several components:

- ▶ Monitor I, Monitor II, Monitor III
- ▶ Postprocessor
- ▶ RMF Performance Monitoring
- ▶ Client/Server Reporter
- ▶ Sysplex Data Server
- ▶ Distributed Data Server
- ▶ LDAP back-end

These components complement each other to provide the infrastructure for performance management:

- ▶ Gathering data
- ▶ Reporting data
- ▶ Accessing data across the sysplex

We introduce data gathering and RMF reporting in the next sections.

### 9.4.1 Data gathering

RMF gathers data using three monitors:

- ▶ Short-term data collection with Monitor III
- ▶ Snapshot monitoring with Monitor II
- ▶ Long-term data gathering with Monitor I and Monitor III

The system operator starts all monitors as non-interactive (background) sessions with a variety of options that determine what type of data is collected and where it is stored. The data gathering functions run independently on each system, but each monitor can be started sysplex-wide by one operator command.

### 9.4.2 RMF reporting

RMF has three monitors and a postprocessor for reporting performance statistics:

- ▶ RMF Monitor I produces interval reports that are created at the end of a measurement interval, for example, 30 minutes. You can obtain Monitor I session reports during or at the end of RMF processing, or they can be generated at a later time by the postprocessor.
- ▶ RMF Monitor II is a snapshot reporting tool used to obtain fast information about how specific address spaces or system resources (processor, DASD, volumes, storage) are performing. Monitor II has two modes for reporting on system performance:
  - A Monitor II display session, which can be invoked from either an ISPF dialog or directly with a TSO/E command (RMFMON)
  - A Monitor II background session, which is a non-interactive session to create a report for printing

- ▶ RMF Monitor III runs as a TSO/E session under ISPF and provides sysplex or system performance reports by:
  - Displaying your current system status in real-time mode
  - Showing previously collected data that is still available in either storage buffers or preallocated VSAM data sets

Monitor III offers a wide spectrum of reports for answering questions that arise during the various performance management tasks. All reporting is available within one TSO/E session, so there is no need to log on to different systems in the sysplex to get all performance data.

- ▶ The postprocessor is usually invoked using a batch job, but it can also be run under a TSO/E session. It offers the following types of reports:
  - Interval reports, which reflect a picture of performance for each interval for which the data was gathered
  - Duration reports, which summarize data over longer periods of time with a maximum value of 100 hours
  - Summary, plot, exception, and overview reports

In addition, the postprocessor can create overview records that are the optimal base for further spreadsheet processing on the workstation.

## 9.5 RMF example reports

The following examples were produced using the RMF postprocessor. Note that with any RMF reports, the I/O information is reported from *one* z/OS system only. If you are sharing I/O across multiple systems, you will need to review reports from each of the sharing systems to see the complete I/O activity.

### 9.5.1 DASD Activity report

The DASD Activity report, Example 9-1, is the first starting point for any I/O analysis and consequently one we must focus on in some detail.

Example 9-1 DASD Activity

| D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y |      |                       |        |        |        |                 |            |          |          |                     |      |      |      |         |      |      |     |        |       |     |
|---|------|-----------------------|--------|--------|--------|-----------------|------------|----------|----------|---------------------|------|------|------|---------|------|------|-----|--------|-------|-----|
| z/OS V1R10  |      | SYSTEM ID SC33        |        |        |        | DATE 06/23/2009 |            |          |          | INTERVAL 10.00.036  |      |      |      | PAGE 43 |      |      |     |        |       |     |
|   |      | RPT VERSION V1R10 RMF |        |        |        | TIME 15.39.37   |            |          |          | CYCLE 1.000 SECONDS |      |      |      |         |      |      |     |        |       |     |
| -   |      | TOTAL SAMPLES =       | 600    | IODF = | 07     | CR-DATE:        | 06/19/2009 | CR-TIME: | 10.00.10 | ACT: ACTIVATE       |      |      |      |         |      |      |     |        |       |     |
| STORAGE   | DEV  | DEVICE                | NUMBER | VOLUME | PAV    | LCU             | ACTIVITY   | RESP     | IOSQ     | CMR                 | DB   | PEND | DISC | CONN    | %    | %    | %   | AVG    | %     | %   |
| GROUP   | NUM  | TYPE                  | OF     | CYL    | SERIAL | RATE            | TIME       | TIME     | DLY      | DLY                 | TIME | TIME | TIME | CONN    | DEV  | DEV  | DEV | NUMBER | ANY   | MT  |
| ACIRLS  | D308 | 33909                 | 10017  | ACIV01 | 1.0H   | 0028            | 0.008      | .973     | .000     | .000                | .000 | .128 | .000 | .845    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D309 | 33909                 | 10017  | TBIG08 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D30A | 33909                 | 10017  | SAPTL3 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D30D | 33909                 | 10017  | NWD30D | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D30E | 33909                 | 10017  | IDGC00 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D30F | 33909                 | 10017  | ##\$X2 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D310 | 33909                 | 10017  | FKDBU3 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
|   | D311 | 33909                 | 10017  | COMST4 | 1.0H   | 0028            | 448.748    | .250     | .000     | .018                | .000 | .144 | .002 | .105    | 4.63 | 4.69 | 0.0 | 1.5    | 100.0 | 0.0 |
|   | ...  |                       |        |        |        |                 |            |          |          |                     |      |      |      |         |      |      |     |        |       |     |
|   | D320 | 33909                 | 10017  | TBIG26 | 1.0H   | 0028            | 0.000      | .000     | .000     | .000                | .000 | .000 | .000 | .000    | 0.00 | 0.00 | 0.0 | 0.0    | 100.0 | 0.0 |
| 0   |      |                       |        | LCU    |        | 0028            | 448.780    | .250     | .000     | .018                | .000 | .144 | .002 | .105    | 0.05 | 0.05 | 0.0 | 1.5    | 100.0 | 0.0 |

The fields of most interest are:

- ▶ **DEVICE ACTIVITY RATE:** The rate, in operations per second, at which start subchannel (SSCH) instructions to the device completed successfully.
- ▶ **AVG RESP TIME:** This is the average response time, in milliseconds, for an I/O to this device or LCU. This can be considered as the overall measurement that reflects the health of the device. Response time is the sum of IOSQ, connect, disconnect, and pending time as described later.
- ▶ **AVG IOSQ TIME:** This is the average time, in milliseconds, that an I/O request must wait on the IOS queue before a SSCH instruction can be issued.
- ▶ **AVG CMR (command response) Delay:** This delay indicates the time between a successful I/O operation being initiated on the channel and the device acknowledging the connection. These delays can be caused by contention in the fabric and at the destination port.
- ▶ **AVG DB (device busy) Delay:** This delay is due to the device being busy because of I/O from another sharing z/OS system.
- ▶ **AVG PEND TIME:** The average number of milliseconds an I/O request (SSCH) must wait in the channel subsystem before being accepted at the device. The value includes the time waiting for an available channel path, control unit, and the delay due to shared DASD contention.
- ▶ **AVG DISC TIME:** This reflects the time when the device was in use but not transferring data.
- ▶ **AVG CONN TIME:** This is the time measured by the channel subsystem during which the device is actually connected to the CPU through the path (channel, control unit, DASD) and transferring data.

Figure 9-3 illustrates when and how IOSQ, Pend, and Connect times appear in a DASD I/O operation when a cache hit occurs.

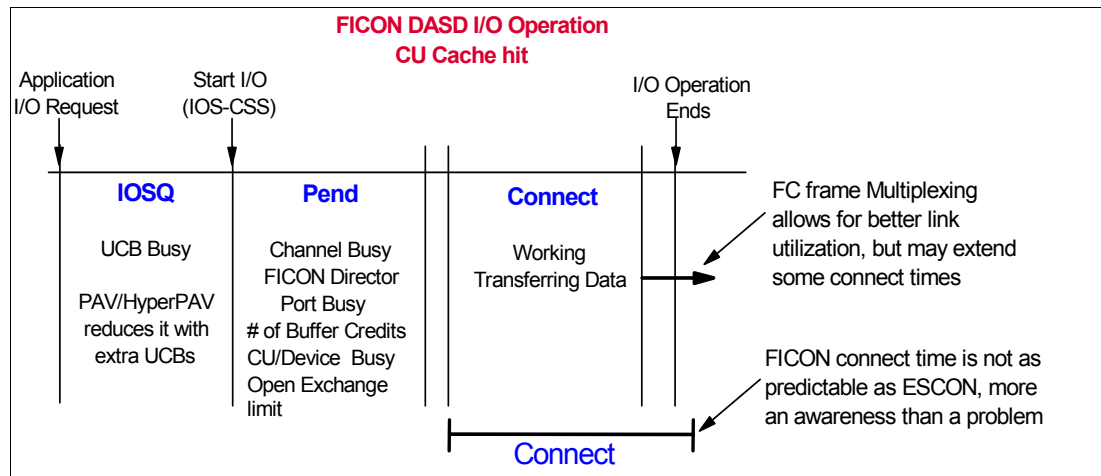


Figure 9-3 FICON DASD I/O operation times with a CACHE hit

Figure 9-4 on page 325 illustrates when and how IOSQ, Pend, Disconnect, and Connect times appear in a DASD I/O operation when a cache misses or an extent conflict has to be resolved by the control unit.

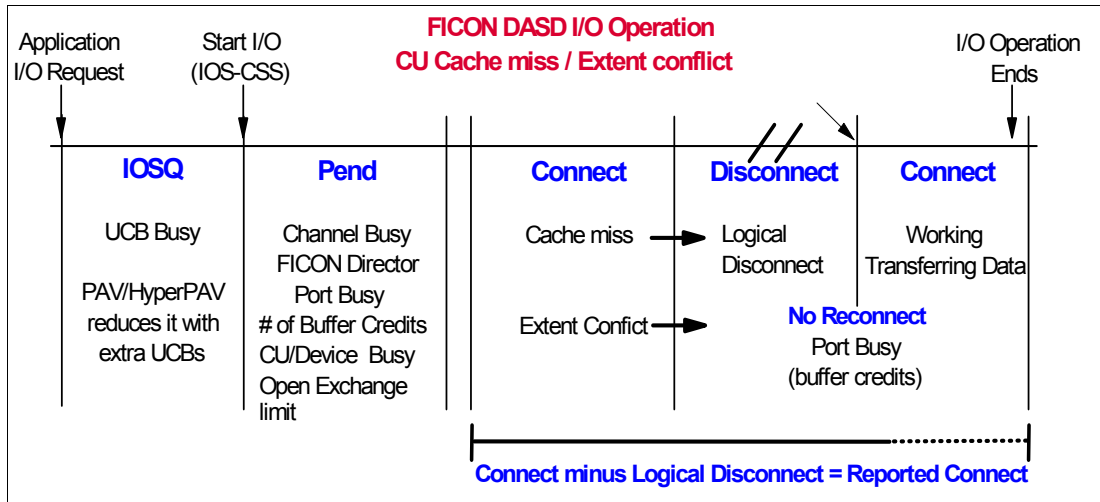


Figure 9-4 FICON DASD I/O operation times with a cache miss or an extent conflict

## 9.5.2 I/O Queueing report

The I/O Queueing report highlights percentages of activity rates and queue lengths for I/O components grouped by Logical Control Unit (LCU). A sample report is shown in Example 9-2.

Example 9-2 I/O Queueing Activity

| z/OS V1R10           |          | I/O QUEUING ACTIVITY |       |                     |                |                   |      |                        |      |                     |       | DATE 06/23/2009 |      |      |      | INTERVAL 10.00.037 |  |  |  |
|----------------------|----------|----------------------|-------|---------------------|----------------|-------------------|------|------------------------|------|---------------------|-------|-----------------|------|------|------|--------------------|--|--|--|
|                      |          | SYSTEM ID SC30       |       | RPT VERSION V1R10   |                | RMF               |      | TIME 15.39.37          |      | CYCLE 1.000 SECONDS |       |                 |      |      |      |                    |  |  |  |
| -TOTAL SAMPLES = 600 |          | IODF = 07            |       | CR-DATE: 06/19/2009 |                | CR-TIME: 10.00.10 |      | ACT: ACTIVATE          |      |                     |       |                 |      |      |      |                    |  |  |  |
| - INITIATIVE QUEUE - |          | IOP UTILIZATION      |       | I/O START           |                | INTERRUPT         |      | % I/O REQUESTS RETRIED |      | RETRIES / SSCH      |       |                 |      |      |      |                    |  |  |  |
| IOP                  | ACTIVITY | AVG Q                | % IOP | I/O RATE            | INTERRUPT RATE | ALL               | CP   | DP                     | CU   | DV                  | ALL   | CP              | DP   | CU   | DV   |                    |  |  |  |
|                      | RATE     | LNTH                 | BUSY  | RATE                | RATE           |                   | BUSY | BUSY                   | BUSY | BUSY                |       | BUSY            | BUSY | BUSY | BUSY |                    |  |  |  |
| 0 00                 | 1454.535 | 0.00                 | 3.14  | 1454.535            | 1338.536       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 01                   | 1452.405 | 0.00                 | 1.17  | 1452.405            | 1538.245       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 02                   | 1453.399 | 0.00                 | 1.33  | 1453.397            | 2342.267       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 03                   | 1451.525 | 0.00                 | 1.47  | 1451.525            | 2460.060       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 04                   | 1452.490 | 0.00                 | 1.42  | 1452.490            | 2161.580       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 05                   | 1451.754 | 0.00                 | 2.33  | 1451.754            | 1575.881       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| 0 SYS                | 8716.105 | 0.00                 | 1.81  | 8716.105            | 11416.57       | 0.0               | 0.0  | 0.0                    | 0.0  | 0.0                 | 0.00  | 0.00            | 0.00 | 0.00 | 0.00 |                    |  |  |  |
| LCU                  | CU       | DCM                  | GROUP | CHAN                | CHPID          | % DP              | % CU | CUB                    | CMR  | CONTENTION          | DELAY | AVG             | DATA |      |      |                    |  |  |  |
|                      |          | MIN                  | MAX   | DEF                 | TAKEN          | BUSY              | BUSY | DLY                    | DLY  | RATE                | LNTH  | DLY             | WAIT | MAX  | EXCH | CONC               |  |  |  |
| 0 0001               | 2080     |                      |       |                     | 02             | 0.127             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.127             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 0002               | 20A0     |                      |       |                     | 03             | 0.127             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.127             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 0003               | 20C0     |                      |       |                     | 04             | 0.133             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.133             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 0004               | 20E0     |                      |       |                     | 05             | 0.127             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.127             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 0007               | 3080     |                      |       |                     | 1C             | 0.127             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.127             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 0008               | 30A0     |                      |       |                     | 1E             | 0.127             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.127             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.0             |      |      |      |                    |  |  |  |
| 0 000B               | 4B00     |                      |       |                     | 4C             | 0.317             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.317             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.1             |      | 0.00 | 0.00 |                    |  |  |  |
| 0 000C               | 4B08     |                      |       |                     | 4E             | 0.013             | 0.00 | 0.00                   | 0.0  | 0.0                 | 0.00  | 0.00            |      |      |      |                    |  |  |  |
|                      |          |                      |       |                     | *              | 0.013             | 0.00 | 0.00                   | 0.0  | 0.000               | 0.00  | 0.1             |      | 0.00 | 0.00 |                    |  |  |  |

In an IBM PR/SM™ environment (such as our configuration), the report is split into two sections. The top section reports PR/SM system activity. The rest of the report applies to I/O activity for the z/OS system being measured by RMF. The fields we focus on are:

- ▶ PR/SM fields from the I/O Queuing report:
  - IOP: The two digit hexadecimal identifier of the I/O processor (IOP), which is part of the channel subsystem. IOPs refer to the System Assist Processors (SAP) on your server.
  - INITIATIVE QUEUE: This queue reflects the I/O initiation activity for each I/O processor (IOP). Average queue length reflects the average number of entries on the initiative queue for each IOP. The activity rate reflects the assignment of I/O requests.
  - IOP UTILIZATION shows fields to measure the I/Os started and interrupted on each IOP and an overall percentage busy indicator.
  - I/O REQUESTS RETRIED reflects the ratio of retried I/O requests. RETRIES/SSCH (start subchannel) reflects the retried I/O requests initially started. Reasons for each:
    - CP - Channel path busy
    - DP - Director port busy
    - CU - Control unit busy
    - DV - Device busy
- ▶ z/OS System activity fields from the I/O Queuing report:
  - LCU and CU are the logical and physical control unit identifiers. They show the logical assignment to the physical resources.
  - The Dynamic Channel Path Management (DCM) GROUP reports the minimum, maximum, and initially defined number of DCM channels for each LCU in this interval.
  - CHPID TAKEN shows how evenly I/O requests are spread across the available paths to the LCU.
  - The %BUSY fields show the percentage of I/O requests that were deferred during the interval due to director port (DP) or control unit (CU) resource being busy.
  - AVG CUB DLY indicates the time (in milliseconds) that an I/O operation was delayed due to the control unit being busy.
  - AVG CMR DLY indicates the time (in milliseconds) between a successful I/O operation being initiated and the device acknowledging the connection.
  - CONTENTION RATE indicates the rate that the I/O processor places the request on the control unit header for this interval. This occurs when all paths to the subchannel are busy and at least one path to the control unit is busy.
  - DELAY Q LENGTH reflects the average number of delayed requests on the control unit header for this interval.
  - AVQ CSS DLY reflects a time (in milliseconds) that an I/O is delayed from the acceptance of a start on the subchannel to the point that the channel subsystem actually initiates the operation on this LCU.

### 9.5.3 Channel Path Activity report

Example 9-3 on page 327 shows a Channel Path Activity report. The purpose of this report is to reflect channel path activity for the requested interval.

Example 9-3 Channel Path Activity

| CHANNEL PATH ACTIVITY |      |                       |     |      |                 |      |      |                     |      |            |                  |                     |       |                 |          |       |
|-----------------------|------|-----------------------|-----|------|-----------------|------|------|---------------------|------|------------|------------------|---------------------|-------|-----------------|----------|-------|
| z/OS V1R10            |      | SYSTEM ID SC30        |     |      | DATE 06/23/2009 |      |      | INTERVAL 10.00.037  |      |            | PAGE 2           |                     |       |                 |          |       |
|                       |      | RPT VERSION V1R10 RMF |     |      | TIME 15.39.37   |      |      | CYCLE 1.000 SECONDS |      |            |                  |                     |       |                 |          |       |
| IODF = 07             |      |                       |     |      |                 |      |      |                     |      | MODE: LPAR |                  | CPMF: EXTENDED MODE |       |                 | CSSID: 2 |       |
| CHANNEL PATH          |      | UTILIZATION(%)        |     |      | READ(MB/SEC)    |      |      | WRITE(MB/SEC)       |      |            | FICON OPERATIONS |                     |       | ZHPF OPERATIONS |          |       |
| ID                    | TYPE | G                     | SHR | PART | TOTAL           | BUS  | PART | TOTAL               | PART | TOTAL      | RATE             | ACTIVE              | DEFER | RATE            | ACTIVE   | DEFER |
| 51                    | FC_S |                       |     |      | OFFLINE         |      |      |                     |      |            |                  |                     |       |                 |          |       |
| 52                    | FC_S | 4                     | Y   | 9.39 | 15.66           | 1.30 | 0.36 | 0.75                | 1.96 | 3.23       | 1176.8           | 1.5                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 53                    | FC_S | 4                     | Y   | 0.01 | 0.29            | 0.05 | 0.01 | 0.15                | 0.00 | 0.08       | 18.1             | 1.1                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 54                    | FC_S |                       |     |      | OFFLINE         |      |      |                     |      |            |                  |                     |       |                 |          |       |
| 55                    | FC_S |                       |     |      | OFFLINE         |      |      |                     |      |            |                  |                     |       |                 |          |       |
| 56                    | FC_S | 4                     | Y   | 9.37 | 15.63           | 1.30 | 0.35 | 0.75                | 1.97 | 3.24       | 1176.3           | 1.5                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 57                    | FC_S | 4                     | Y   | 0.01 | 0.23            | 0.05 | 0.01 | 0.16                | 0.01 | 0.09       | 16.0             | 1.1                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 58                    | FC_S | 4                     | Y   | 2.58 | 3.45            | 0.25 | 0.18 | 0.71                | 0.00 | 0.16       | 219.3            | 1.0                 | 0.0   | 90.6            | 1.0      | 0.0   |
| 59                    | FC_S | 4                     | Y   | 2.46 | 4.09            | 0.33 | 0.17 | 1.14                | 0.00 | 0.05       | 221.0            | 1.1                 | 0.0   | 86.8            | 1.0      | 0.0   |
| 5A                    | FC_S | 4                     | Y   | 9.41 | 15.67           | 1.30 | 0.35 | 0.76                | 1.96 | 3.22       | 1178.5           | 1.5                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 5B                    | FC_S | 4                     | Y   | 0.01 | 0.40            | 0.07 | 0.01 | 0.20                | 0.01 | 0.12       | 21.3             | 1.1                 | 0.0   | 0.0             | 0.0      | 0.0   |
| 5C                    | FC_S | 4                     | Y   | 1.77 | 2.51            | 0.20 | 0.13 | 0.58                | 0.00 | 0.14       | 153.4            | 1.0                 | 0.0   | 67.7            | 1.1      | 0.0   |
| 5D                    | FC_S | 4                     | Y   | 1.65 | 3.28            | 0.30 | 0.13 | 1.15                | 0.00 | 0.06       | 157.1            | 1.1                 | 0.0   | 65.4            | 1.0      | 0.0   |
| 5E                    | FC_S | 4                     | Y   | 9.39 | 15.66           | 1.30 | 0.36 | 0.77                | 1.96 | 3.23       | 1178.4           | 1.5                 | 0.0   | 0.0             | 0.0      | 0.0   |

The fields of interest are:

► CHANNEL PATH:

- ID is the CHPID number.
- TYPE is the type of the channel. FC\_S indicates FICON switched channel.
- G reflects the generation type (4 indicates a 4 Gbps channel - FICON).
- SHR indicates whether the channel is shared between one or more partitions.

► UTILIZATION:

- PART reflects the utilization for this partition for this interval.
- TOTAL reflects the utilization from the entire processor for this interval.
- BUS reflects the percentage of cycles the bus was found to be busy for this channel in relation to the potential limit.

► READ(MB/SEC) and WRITE(MB/SEC):

- PART shows the data transfer rate in MBps for this channel from this partition to the control unit in this interval.
- TOTAL shows the data transfer rate in MBps for this channel from the entire processor to the control unit for this interval.

**Note:** On a machine running in LPAR mode, but with only one LPAR defined, the PART columns of this report reflect a zero (0) value for the READ, WRITE, and UTILIZATION displays for FICON channels.

► FICON and zHPF Operations (Physical Channel):

- RATE is the number of native FICON or zHPF operations per second at the physical channel level.
- ACTIVE is the average number of native FICON or zHPF operations that are concurrently active. Often referred to as the number of open exchanges.
- DEFER is the number of deferred native FICON or zHPF operations per second. This is the number of operations that cannot be initiated by the channel because of lack of available resources.

## 9.5.4 FICON Director Activity report

The FICON Director Activity report provides configuration information of connected FICON directors and details pertaining to the connectivity and activity of each director port. The measurements provided for a port in the report do not only comprise the I/O for the system on which the report is taken, but include all I/O that is directed through this port, regardless of which LPAR requests the I/O.

Example 9-4 shows a report for one director and its ports.

Example 9-4 FICON Director Activity

---

```

1                               F I C O N   D I R E C T O R   A C T I V I T Y                               PAGE    1
                                z/OS V1R10                SYSTEM ID SC30                DATE 06/23/2009                INTERVAL 10.00.037
                                RPT VERSION VIR10 RMF        TIME 15.39.37                CYCLE 1.000 SECONDS
- IODF = 07  CR-DATE: 06/19/2009  CR-TIME: 10.00.10  ACT: ACTIVATE
0 SWITCH DEVICE: 0062  SWITCH ID: 62  TYPE: 006064  MODEL: 001  MAN: MCD  PLANT: 01  SERIAL: 0000000119D2
0 PORT  -CONNECTION-  AVG FRAME  AVG FRAME SIZE  PORT BANDWIDTH (MB/SEC)  ERROR
  ADDR  UNIT  ID  PACING  READ  WRITE  -- READ --  -- WRITE --  COUNT
      CU  8400
      CU  8200
      CU  8000
0D SWITCH  ----  0  56  56  0.00  0.00  0
0E CHP  86  0  72  279  0.00  0.00  0
0F CHP  54  0  388  965  0.05  0.17  0
10 CU  ----  0  796  329  0.02  0.00  0
11 CU  ----  0  683  459  0.02  0.01  0
12 CHP  87  0  71  107  0.00  0.00  0
13 CU  ----  0  94  40  0.00  0.00  0
14 CU  ----  0  973  284  0.04  0.01  0
15 CU  ----  0  747  397  0.14  0.07  0
16 CHP  8C  0  72  449  0.00  0.01  0
17 CHP  55  0  397  934  0.11  0.27  0
18 SWITCH  ----  0  106  76  0.00  0.00  0
19 CU  ----  0  455  686  0.07  0.17  0
1A CHP  8D  0  68  171  0.00  0.00  0
1B CHP  45  0  0  0  0.00  0.00  0
1C CU  ----  0  957  111  0.10  0.01  0
1D -----  ----  0  0  0  0.00  0.00  0
1E CHP  8E  0  72  169  0.00  0.00  0
1F CHP  95  0  0  0  0.00  0.00  0
20 CHP  42  0  0  0  0.00  0.00  0
21 SWITCH  ----  0  56  58  0.00  0.00  0

```

---

Fields of significance are:

- ▶ **AVG FRAME PACING:** This is the average time (in microseconds) that a frame had to wait before it can be transmitted due to no buffer credits being available.
- ▶ **AVG FRAME SIZE READ/WRITE:** The average frame size (in bytes) used to transmit and receive data during this interval.
- ▶ **PORT BANDWIDTH READ/WRITE:** The rate (in MBps) of data transmitted and received during the interval.
- ▶ **ERROR COUNT:** The number of errors that were encountered during the interval.

## 9.5.5 Enterprise Disk Systems Link Statistics report

The Enterprise Disk Systems report provides measurements about the activities of an enterprise disk system. RMF monitors the activity on a disk system independently from the source of the activity. Activity can be caused by the z/OS system on which RMF is running or from any other system using the disk system. You can use the data contained in this report for checking your current disk configuration, for bottleneck analysis, and for capacity planning.



If the RMF Monitor I data gatherer is set up to collect all available data, the report contains the following three sections:

- ▶ ESS Link Statistics
- ▶ ESS Extent Pool Statistics
- ▶ ESS Rank Statistics

Extent Pool and Rank Statistics are available for the IBM System Storage DS® disk system family. They provide capacity and performance information about allocated disk space in the disk system and activity statistics about read and write operations on the disk arrays. Use these reports to detect and resolve performance problems and for disk capacity planning.

The ESS Link Statistics section must be used for analysis of the external link usage and for capacity planning of peer-to-peer remote copy (PPRC) links. For each disk system host adapter port (SAID), this report contains statistics by link type. Example 9-5 shows the activity on adapters of one disk system.

*Example 9-5 ESS Link Statistics*

---

```

1                               E S S L I N K S T A T I S T I C S                               PAGE 2
                                z/OS V1R10          SYSTEM ID SC31          DATE 05/13/2009          INTERVAL 10.00.237
                                RPT VERSION V1R10 RMF          TIME 21.49.36          CYCLE 1.000 SECONDS
OSERIAL NUMBER 00000BALB1    TYPE-MODEL 002107-922  CDATE 05/13/2009  CTIME 21.49.37  CINT 10.00
0-----ADAPTER-----  --LINK TYPE--  BYTES  BYTES  OPERATIONS  RESP TIME  I/O
SAID  TYPE              /SEC    /OPERATION  /SEC          /OPERATION  INTENSITY
00201 FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00202 FIBRE 2Gb        ECKD READ      1.7M      5.4K      315.1      0.1      19.1
                                ECKD WRITE     142.4K    1.8K      78.5       0.0       3.7
                                -----
                                22.8
00203 FIBRE 2Gb        ECKD READ      899.4K    5.2K      174.0      0.1      9.8
                                ECKD WRITE     83.7K     2.3K      37.1       0.1       2.1
                                -----
                                11.9
00230 FIBRE 2Gb        ECKD READ      2.4M     618.8     3808.8     0.0     65.1
                                ECKD WRITE     661.5K    11.4K     58.1       0.1       7.6
                                -----
                                72.7
00231 FIBRE 2Gb        ECKD READ      2.1M     553.0     3799.4     0.0     63.7
                                ECKD WRITE     633.1K    12.9K     49.2       0.2       8.7
                                -----
                                72.3
00232 FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00233 FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00240 FIBRE 2Gb        ECKD READ      168.4K    9.6K      17.5       0.1       1.0
                                ECKD WRITE     857.2K    20.3K     42.3       0.2       6.7
                                -----
                                7.6
00241 FIBRE 2Gb        ECKD READ      45.4K     3.8K      12.0       0.0       0.4
                                ECKD WRITE     9.2K      1.8K      5.0        0.1       0.5
                                -----
                                0.9
00242 FIBRE 4Gb        ECKD READ      5.5K      2.6K      2.1        0.0       0.1
                                ECKD WRITE     655.4     1.9K      0.3        0.1       0.0
                                -----
                                0.1
00243 Undefined        NO DATA TO REPORT OR ZERO
00300 FIBRE 2Gb        ECKD READ      2.6M     676.2     3843.2     0.0     67.8
                                ECKD WRITE     745.8K    10.2K     72.9       0.1       9.2
                                -----
                                77.0

```

---

The fields are:

- ▶ ADAPTER:
  - SAID reflects the System Adapter Identifier address of the host adapter port.
  - TYPE reflects the characteristics of the connection (2 Gb or 4 Gb fiber in our example).

- ▶ LINK TYPE will be one of the following:
  - IBM ECKD™ read and write: Designates extended count key data I/O
  - SCSI read and write: Designates small computer system interface I/O
  - PPRC send and receive: Designates peer-to-peer remote copy traffic
- ▶ BYTES/SEC reflects the average number of bytes transferred for all operations on this link.
- ▶ BYTES/OPERATION reflects the average number of bytes transferred for each individual operation.
- ▶ OPERATIONS/SEC reflects the average number of operations per second.
- ▶ RESP TIME/OPERATION reflects the average response time (in milliseconds).
- ▶ I/O INTENSITY is measured in milliseconds/second. It reflects the duration of the interval for which the adapter was active. A value of 1000 indicates that the link was busy for the entire time period. A value greater than 1000 would indicate that concurrent operations had been active.

## 9.5.6 General performance guidelines

The following guidelines are basic practices that can be used to improve the performance of your FICON devices. Be aware that their applicability can vary based on vendor specifications and model type.

### **DASD**

In this section, we discuss basic practices for DASD.

#### ***High disconnect time***

The following points reflect typical approaches to reduce disconnect time:

- ▶ Cache the volume, if it is a suitable cache candidate.
- ▶ Un-cache the volume, if it is a poor cache candidate.
- ▶ Tune cache.
- ▶ Used Dynamic Cache Management (DCME) to dynamically manage cache for you.
- ▶ If necessary, manually tune cache by turning off poor cache users.
- ▶ Review blocksize. Increase where applicable.
- ▶ Add paths (if not already 8-path).
- ▶ Use faster devices. For highest performance use Solid State Disks (SSD).
- ▶ Use disk system functions, such as IBM DS8000 Easy Tier®, which automatically move high activity data to faster devices.
- ▶ Move or copy data sets to other volumes to reduce contention.
- ▶ Use I/O Priority Manager, an optional feature of DS8700 and DS8800 disk systems, which enables more effective storage consolidation and performance management with the ability to align quality of service levels to separate workloads in the system.

#### ***High IOSQ time***

The following points reflect typical approaches to reduce IOSQ time:

- ▶ Implement Parallel Access Volumes, preferably HyperPAV.
- ▶ Decrease disconnect, connect, and PEND times. Decreasing the other response time components will also decrease IOSQ time proportionally.

- ▶ Cache the volume, if it is a suitable cache candidate.
- ▶ Tune cache.
- ▶ Use Dynamic Cache Management (DCME) to manage cache for you.
- ▶ Use I/O priority queuing (IOQ=PRTY in IEAIPSxx).
- ▶ Move or copy data sets to other volumes to reduce contention.
- ▶ Run requests sequentially to reduce contention.

### ***High connect time***

The following points reflect typical approaches to reduce connect time:

- ▶ Cache the volume, if it is a suitable cache candidate.
- ▶ Use indexed VTOCs.
- ▶ In a cached environment, use faster channels.
- ▶ Use PDSEs if you suspect directory search time is a problem.

### ***High PEND time***

The following points reflect typical approaches to reduce PEND time:

- ▶ Change the mix of data on the volume to reduce contention. If you can, identify one data set that is contributing to most of the problem. This can be eligible to be moved to a custom volume or moved into storage.
- ▶ Check channel utilization, especially if not a small number of channels per controller.
- ▶ Exploit multiple allegiance if available on your DASD subsystem.
- ▶ Implement Parallel Access Volumes.
- ▶ Check CMR time and estimate CU host adapter utilization to see if CU is being overloaded. Low channel utilization with high pend times means you might have to add more CU host adapters.

## **9.5.7 Tape**

The following list items reflect typical approaches to improving tape performance. These are basic recommendations. The significant developments in tape technology (such as the IBM Virtual Tape Server) mitigate many of these points, and you must refer to the specific performance recommendations of your chosen solution for more guidance.

- ▶ Review blocksize; increase where applicable.
- ▶ Allocate data sets to DASD. Tape mount management (TMM) is a method that you can use to accomplish this, and thereby significantly reduce tape mount requirements.
- ▶ Increase host buffers for priority jobs. Increase buffers from the default of 5, to as high as 20 (but no more than 20).
- ▶ Use VIO for temporary data sets.
- ▶ Add channels (if not already one channel per control unit function).
- ▶ Reduce mount delay (for example, with Automated Cartridge Loader).
- ▶ Use enhanced-capacity cartridges and tape drives.
- ▶ Restructure for more parallelism (for example, to allow an application to use multiple drives).

## 9.6 IBM Network Advisor performance monitor

IBM Network Advisor performance monitoring provides details about the quantity of traffic and of errors a specific port or device generates on the fabric over a specific time frame. You can use performance monitor to indicate the devices that create the most traffic and to identify the ports that are most congested.

The IBM Network Advisor performance monitor allows you to monitor your SAN using the following methods:

- ▶ Gather and display real-time performance data for selected ports.
- ▶ Collect and display historical performance data for ports in selected fabrics or in the entire SAN.
- ▶ Support end-to-end monitors for monitoring data flow between two ports in a fabric for both real-time and historical performance data.
- ▶ Display port-to-port connections which are using the most bandwidth on the selected director or one of the F-ports on the director with a feature called Top Talkers.

In this section, we describe how to set up and use the monitors. The screen images are based on IBM Network Advisor v11.1. Refer to *IBM System Storage Network Advisor SAN User Manual*, GA32-0940, for more details.

### 9.6.1 Real-Time Performance Reporting

With the real-time reporter, you can monitor the current performance of selected director ports. The data is displayed in graphical format and is continuously updated on your screen. You can customize the graphs and have the option to export the data in a file.

To set up real-time performance reporting for selected ports:

1. Select a fabric or director on the IBM Network Advisor Product List or Topology Map view, and then select **Monitor** → **Performance** → **Real-Time Graph** from the menu bar. Or right-click a fabric or director, and select **Performance** → **Real-Time Graph** (Figure 9-5 on page 333).

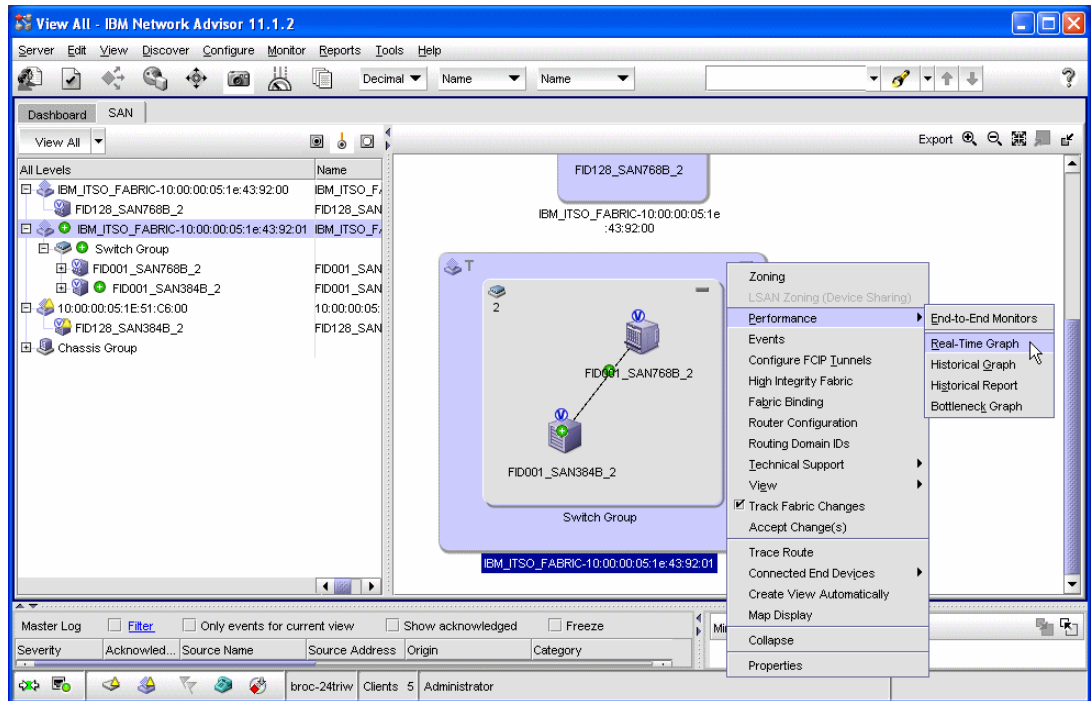


Figure 9-5 Real-Time Monitor set up

2. The Real Time Port Selector window opens. Select from the Show drop-down menu the port type you want to monitor (for example, Device Ports or ISL Ports). Ports of the selected type are included in the Available table on the left. Expand the list to find the ports of interest. To expand the whole list, right-click anywhere in the Available table, and select **Expand All**.
3. Select one or more ports on the Available table, and click the right arrow in the middle to add the ports on the monitored port table on the right, as shown in Figure 9-6.

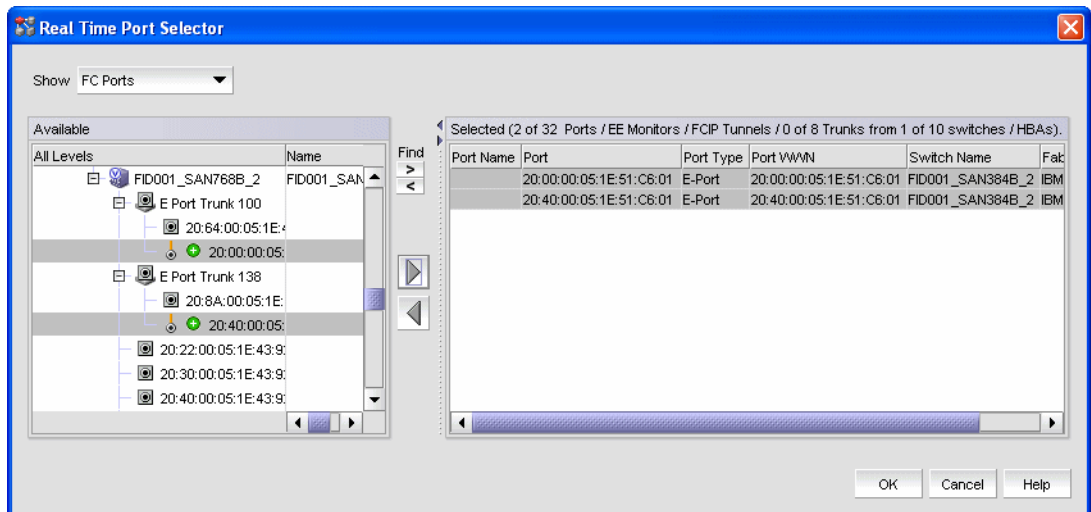


Figure 9-6 Real-Time Port Selector

4. Click **OK** to save the changes. The Real Time Performance Graphs window displays (Figure 9-7 on page 334).

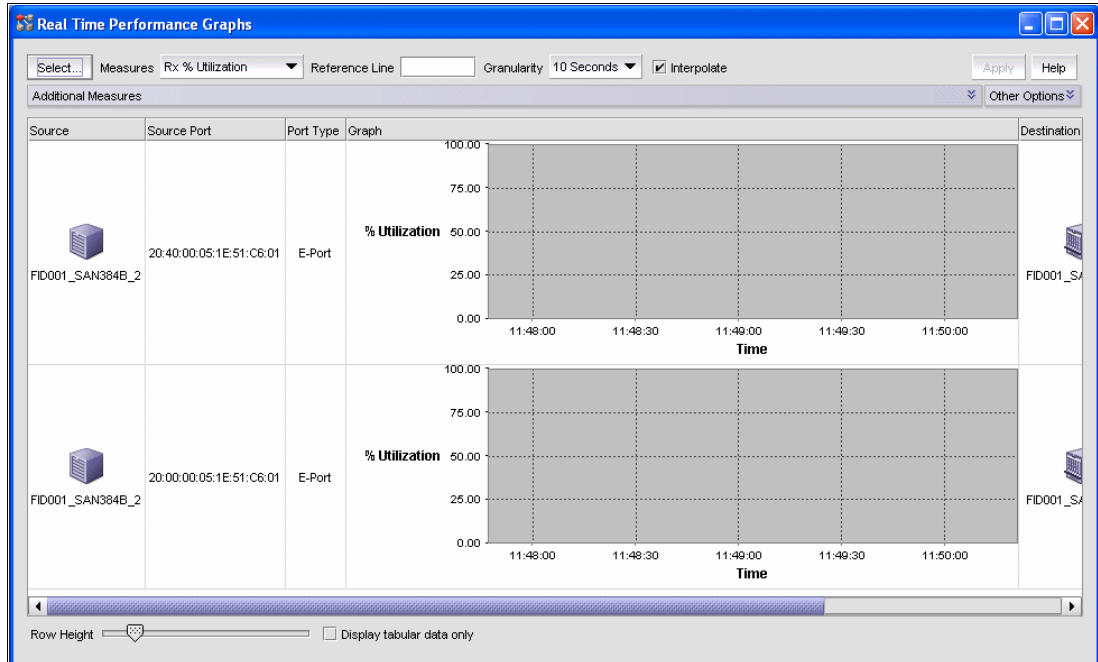


Figure 9-7 Real-Time Performance Graphs

The real-time monitor is now set up and can be used. You will get the real-time performance data of the selected objects.

### Customizing real-time performance graphs

The Real Time Performance Graphs panel contains several buttons which you can use to filter and modify the format of data displayed on the panel.

- ▶ Use the Select button to add ports to or delete ports from the panel.
- ▶ Use the Measures drop-down menu to select which metric is shown on the graph. The available options include:
  - Transmit and receive Utilization %
  - Transmit and receive MBps rate
  - Different error rates, such as CRC errors, signal and sync losses, link failures, and so on.
- ▶ Use the Additional Measures drop-down menu to show multiple metrics on the graph.
- ▶ Enter a value in the Reference Line field to draw a reference line for transmit and receive utilization measurements.
- ▶ Use the Granularity drop-down to set the graph refresh interval between 10 seconds and one minute.
- ▶ You can adjust the graph size or choose to display tabular data only.

After you modify the options, click **Apply** to activate the changes.

### Exporting and printing real-time data

To export real-time performance data to a file for further analysis, right-click anywhere on the Real Time Performance Graphs panel, and select **Export Table**. This exports data for all the ports currently included on the panel. You also have the option to export data for a single port only by selecting **Export Row**, or to print the graph by selecting **Print**.

## Clearing port counters

To reset all port statistics to zero on a selected director, right-click the director on the IBM Network Advisor Product List or Topology Map view, and select **Performance** → **Clear Counters**.

## 9.6.2 Historical performance reporting

Historical performance reporting allows you to create historical performance graphs and tabular reports that you can customize to your needs. You also have the option to export data and reports in a file.

Historical performance reporting requires that performance data is first collected before reports can be created. After the data is collected over a period of time, you can create reports for any interval during which data collection was enabled. Performance data collection must be enabled constantly to receive the necessary historical data required for a meaningful report.

### Enabling historical performance data collection

You can enable historical performance data collection for some or all fabrics in the SAN.

To enable data collection for all fabrics, select **Monitor** → **Performance** → **Historical Data Collection** → **Enable SAN Wide** from the Network Advisor menu bar, as shown in Figure 9-8.

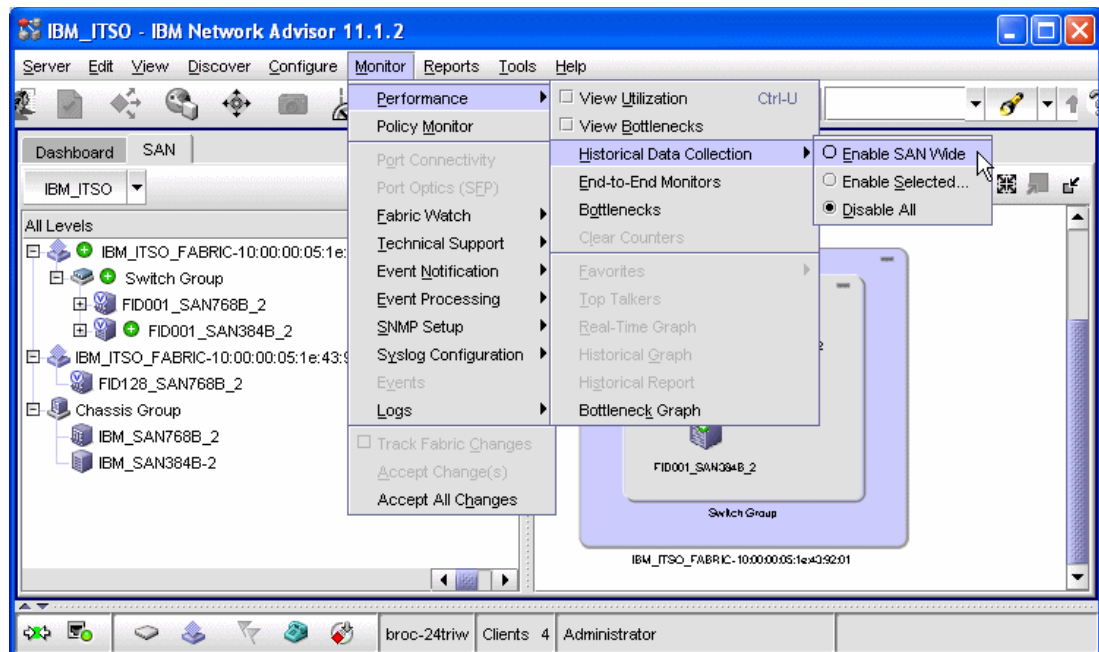


Figure 9-8 Enable SAN wide performance data collection

To enable historical data collection for selected fabrics:

1. Select **Monitor** → **Performance** → **Historical Data Collection** → **Enable Selected** from the Network Advisor menu bar. The Historical Data Collection window opens (Figure 9-9 on page 336).

2. Select one or more fabrics in the Available table on the left, and then click the right arrow in the middle to add them on the Selected table, as shown in Figure 9-9 on page 336. Click **OK** to finish.

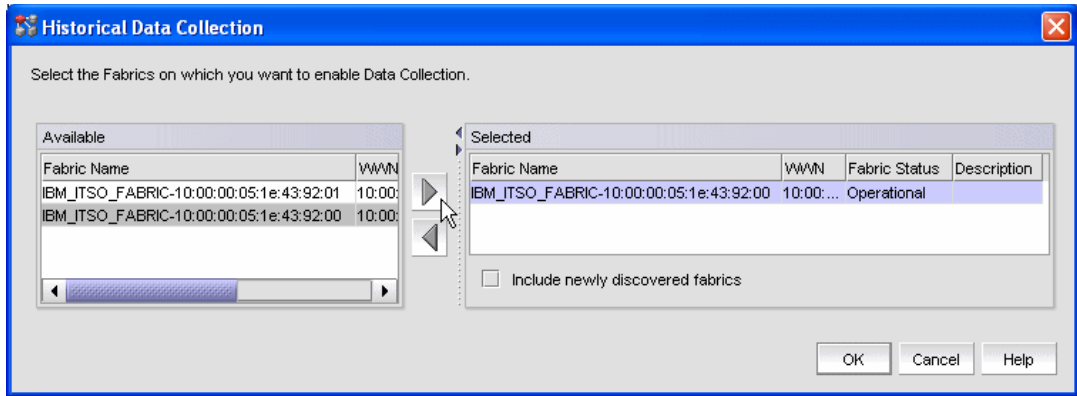


Figure 9-9 Select fabrics

After data collection is enabled, port performance measurements are collected and stored in the IBM Network Advisor server database until data collection is disabled. To disable historical data collection, select **Monitor** → **Performance** → **Historical Data Collection** → **Disable All** from the Network Advisor menu bar.

The time for which historical performance data measurements are stored in the IBM Network Advisor server database is defined on the Server Management Console (SMC). The SMC is an automatically installed, stand-alone application for managing the IBM Network Advisor application server.

There are two predefined aging options available from which you can select:

- ▶ Option 1: 2 years data with the following samples (total 1246 samples):
  - Keep data with 5 minute granularity for last 1 day (288 samples)
  - Keep data with 30 minute granularity for last 3 days (144 samples)
  - Keep data with 2 hour granularity for last 7 days (84 samples)
  - Keep data with 1 day granularity for last 2 years (730 samples)
- ▶ Option 2: 2 years data with the following samples (total 3034 samples):
  - Keep data with 5 minute granularity for last 8 day (2304 samples)
  - Keep data with 1 day granularity for last 2 years (730 samples)

Use the Server Management Console to view and change the aging option. To start the Server Management Console on a Windows server:

1. Logon to the server where IBM Network Advisor server is running, and select **Start** → **Programs** → **IBM Network Advisor** → **Server Management Console**.
2. Open the Performance Data Aging tab (Figure 9-10 on page 337).





Figure 9-10 Performance Data Aging tab

3. Click **Apply** to change the aging option. Note that changing the aging option requires IBM Network Advisor server restart. To do this, start the SMC, open the **Services** tab, and click **Restart** to restart all services.

**Note:** If you change from Option 1 to Option 2, you will lose the existing 30 minute and two hour performance data samples. If you change from Option 2 to Option 1, you will lose the five minute performance data samples, which are older than one day.

### Generating historical performance graphs

To generate a historical performance graph, select a fabric or director on the IBM Network Advisor Product List or Topology Map view, then select **Monitor** → **Performance** → **Historical Graph** from the menu bar. Alternately, right-click a fabric or director, and select **Performance** → **Historical Graph** (Figure 9-11 on page 338).

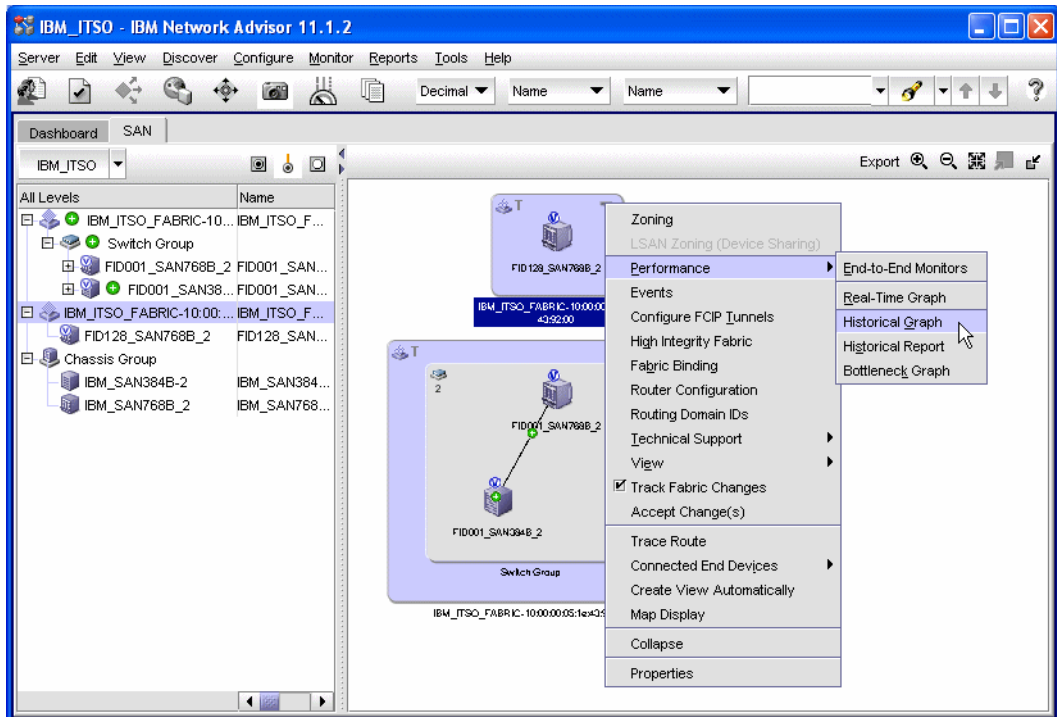


Figure 9-11 Create historical graph

The Historical Performance Graph window shown in Figure 9-12 opens, displaying a graph using the same predefined format that was used the last time a graph was generated.

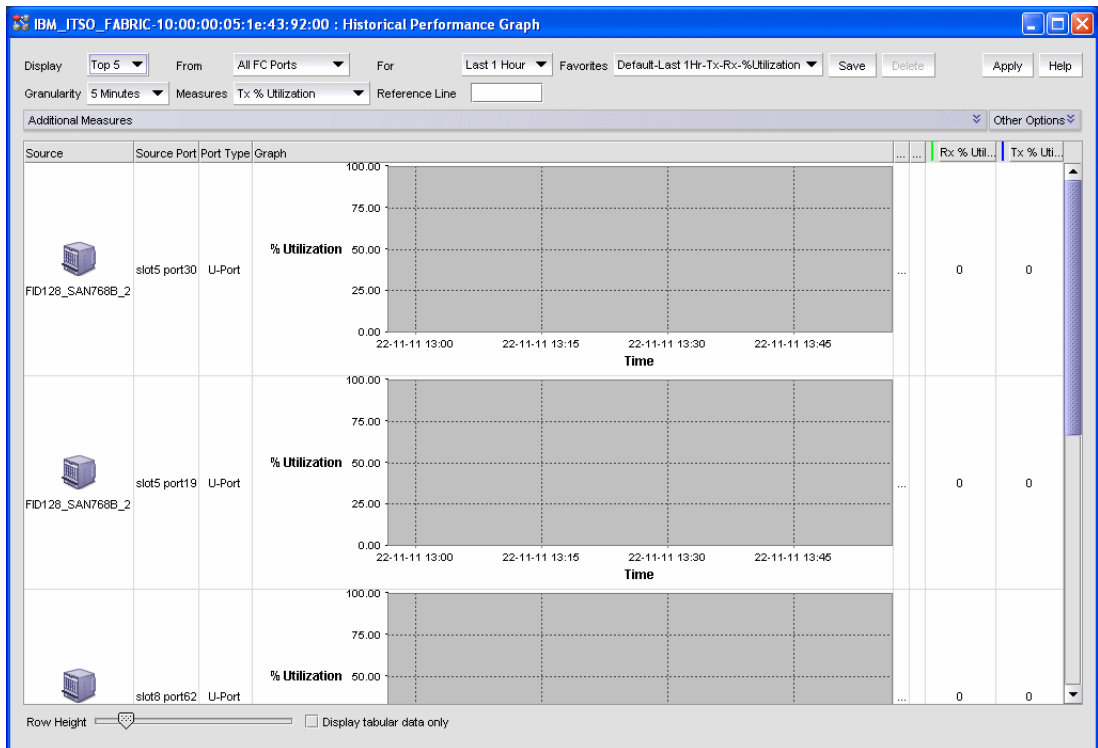


Figure 9-12 Historical Performance Graph panel

## Customizing historical performance graphs

The Historical Performance Graph window contains several buttons that you can use to filter and modify the format of data displayed on the panel:

- ▶ Use the Favorites drop-down menu to select a report configuration from a list of default or custom-saved configurations. The default configuration entries are:
  - Last 1 Hour Tx Rx %Utilization
  - Last 1 Hour Tx MBps
  - Last 1 Day Tx MBps
  - Last 1 Hour Rx MBps
  - Last 1 Day Rx MBps

You can add your own customized configuration in the list.

- ▶ Use the Display drop-down menu to set the number of ports displayed on the panel. Maximum is 20 ports. The most active ports are displayed.
- ▶ Use the From drop-down menu to select which type of ports are displayed. The options are, among others:
  - All FC Ports
  - Device Ports only
  - ISL Ports only
  - Custom, which lets you select the ports which are included in the report
- ▶ Use the For drop-down menu to select the period for which the graphs are displayed. The options are:
  - Last 1 Hour
  - Last 1 Day
  - Last 1 Week
  - Custom, which lets you select the start and end reporting times
- ▶ Use the Granularity drop-down to set the interval at which you want to gather performance data. The minimum is 5 minutes and maximum 1 day.
- ▶ Use the Measures drop-down menu to select the main metric which is shown on the graph. The available metrics are:
  - Transmit and receive Utilization %
  - Transmit and receive MBps rate
  - Different error rates, such as CRC errors, signal and sync losses, link failures etc.
- ▶ Enter a value in the Reference Line field to draw a reference line for the utilization measurements.
- ▶ Use the Additional Measures drop-down menu to show multiple metrics on the graph.
- ▶ You can adjust the graph size or choose to display tabular data only.

After you modify the options, click **Apply** to activate the changes and see the modified graph. You also must click **Apply** to see the most current data because the historical performance graph is not refreshed automatically.

Click **Save** to save your selections as a new custom configuration on the Favorites list. This allows you to easily generate the same kind of graph later. Each time the historical performance graph generator is launched, the first graph is generated using the same Favorites configuration that was used the last time to generate a graph.

Click **Delete** to delete a custom-saved configuration from the Favorites list.

## Exporting and printing historical data

To export historical performance data to a file for further analysis, right-click anywhere on the Historical Performance Graphs panel, and select **Export Table**. This exports data for all of the ports currently included on the panel. You also have the option to export data for a single port only by selecting **Export Row**, or to print the graph by selecting **Print**.

## Generating historical performance reports

Historical performance reports are pre formatted tabular reports that are created using the available historical data. The reports are saved on the IBM Network Advisor server so that you can view them later.

To generate a historical performance report, select a fabric or director on the IBM Network Advisor Product List or Topology Map view, and then select **Monitor** → **Performance** → **Historical Report** from the menu bar. Alternately, you can right-click a fabric or director and select **Performance** → **Historical Report**, as shown in Figure 9-13. A report is generated and automatically displays in the View Reports window. An example report is shown in Figure 9-15 on page 341.

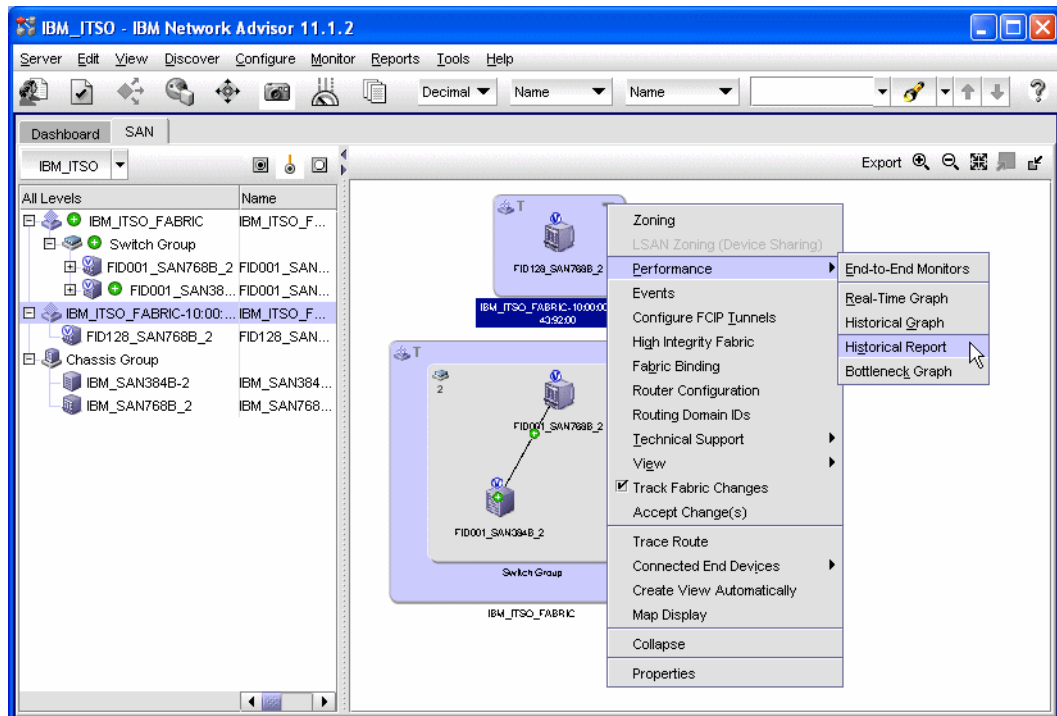


Figure 9-13 Create historical report

When you launch the historical performance report generator, the first report is generated using the same format that was used when the previous historical report was generated. If you want to create a report in another format, for example with different metrics or for a different time period, use the Historical Performance Table window (Figure 9-14 on page 341), which automatically opens when you launch the historical performance report generator.




Figure 9-14 Historical Performance Table

The options and buttons on the Historical Performance Table window are identical to those on the Historical Performance Graph window in Figure 9-12 on page 338. Refer to “Customizing historical performance graphs” on page 339 for instructions on how to use the buttons.

After you change the options, click **Apply** to generate a new report using the modified options. Each time you click **Apply**, a new report is generated and saved even if none of the options was changed.

## Viewing reports

The View Reports window, Figure 9-15, automatically opens when you generate a new report. If you want to view existing reports without generating a new one, open the View Reports window by selecting **Reports** → **View** from the IBM Network Advisor menu bar.

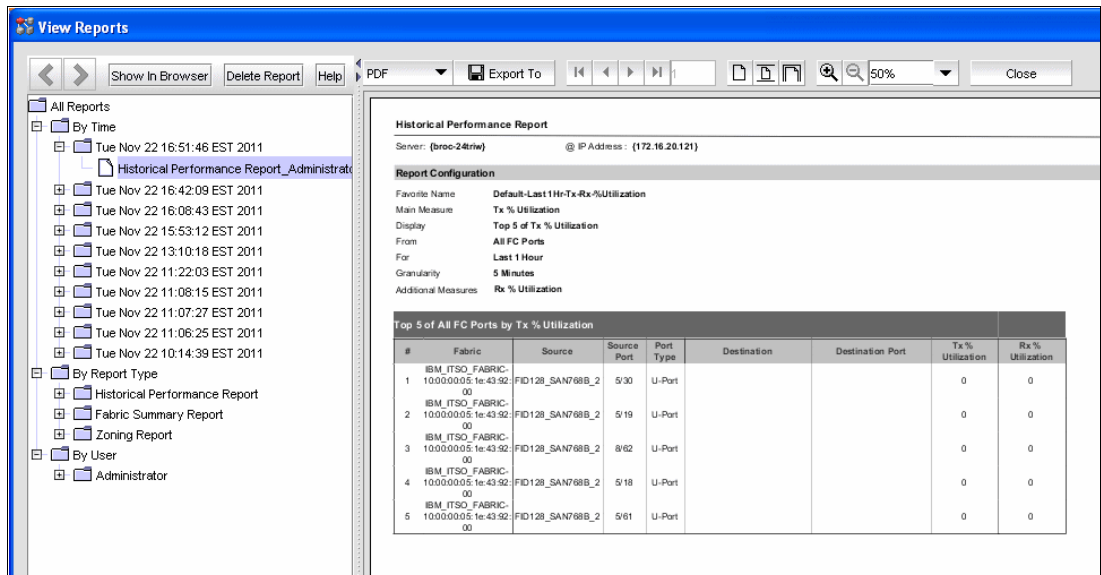


Figure 9-15 View Reports window

The All Reports list on the left shows all reports that are currently available. This list includes not only historical performance reports, but also other types of reports that you or other users generated. You can select any report in the list for viewing by clicking it.

There are several action buttons available on the View Reports panel:

- ▶ Click **Show in Browser** to view the selected report in your default Web browser. You can print the report using the Web browser’s print function. The View Reports window itself does not include print support.
- ▶ Click **Delete Report** to delete the selected report.

- ▶ To save the selected report in a file, first select a format from the drop-down menu (PDF, HTML, or XML), and then click **Export To**.
- ▶ Click **Close** to close the View Reports panel.

### 9.6.3 End-to-End Monitoring

IBM System Storage b-type switches support end-to-end monitoring, which lets you monitor data flow between two ports in a fabric. You can view both real-time and historical performance data for the selected port pair.

**Note:** Either the initiator director or the target director must have an Advanced Performance Monitor license configured to create an end-to-end monitor.

**Note:** A Top Talker and an end-to-end monitor cannot be configured on the same fabric. You must delete the Top Talker monitor before you configure the end-to-end monitor.

To set up an End-to-End Monitor, follow these steps:

1. Select **Monitor** → **Performance** → **End-to-End Monitors** from the IBM Network Advisor menu bar. Alternately, right-click a fabric, and select **Performance** → **End-to-End Monitors**.
2. In the Set End-to-End Monitors panel, Figure 9-16 on page 343, first select the fabric whose ports you want to monitor from the **Fabric** drop-down menu. Select an initiator port (from the upper window on left) and a target port (from the lower window on left). To add the two ports as a pair on the Monitored Pairs table, click the right arrow in the middle. The system automatically determines the initiator SID and the target DID identifiers for the pair and displays them in the Monitored Pairs table.

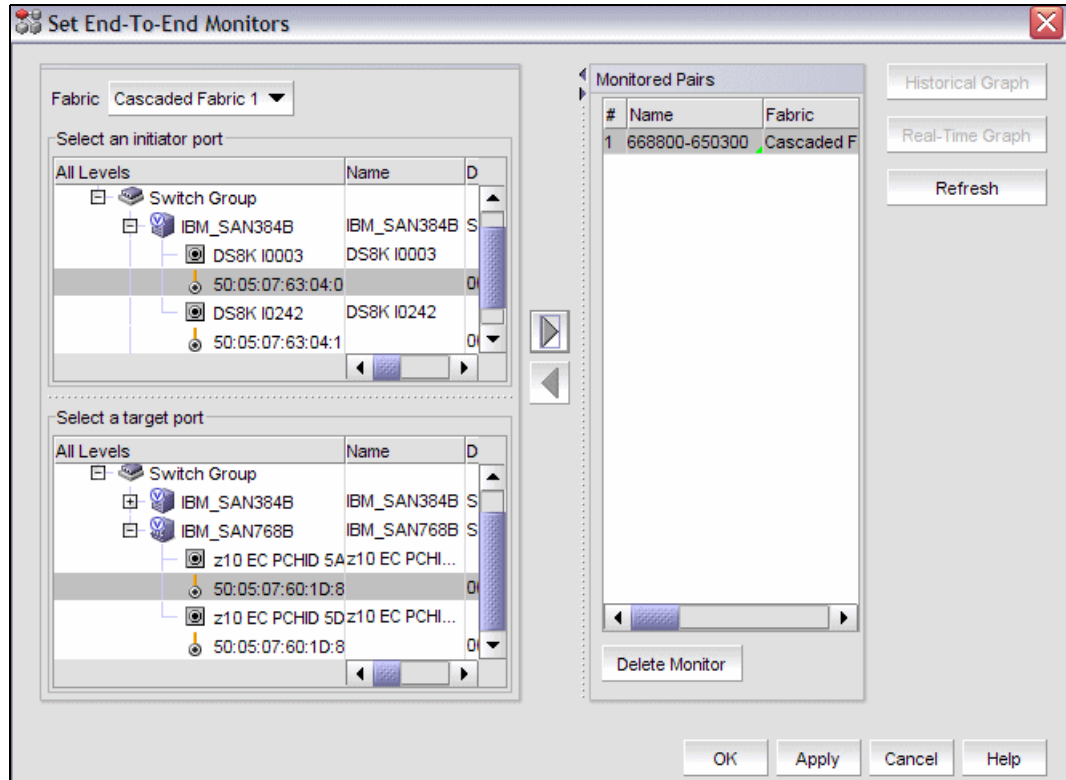


Figure 9-16 End-to-End Monitor setup

Repeat the step to add additional monitored pairs in the table. You can select up to 32 monitored pairs.

3. Click **Apply** to save the changes. After the end-to-end monitored pair is applied to the device, the Status column in the Monitored Pairs table displays 'Enabled'. If the end-to-end monitored pair fails, the Status column in the Monitored Pairs table displays 'Failed:Reason'.

**Note:** If the initiator or target port is part of a logical switch and you move it to another logical switch, the end-to-end monitor fails.

After you create the end-to-end monitored pair, you can view both real-time and historical performance data for the port pair by selecting the pair on the Monitored Pairs table and clicking the **Real-Time Graph** or **Historical Graph** button. Refer to sections 9.6.1, "Real-Time Performance Reporting" on page 332 and 9.6.2, "Historical performance reporting" on page 335 for instructions.

## 9.6.4 Top Talker Monitoring

Use Top Talkers monitor to display the connections that use the most bandwidth on the selected director or port. Top Talkers can be enabled on the director or on one of the F-ports on the director. You can only use Top Talkers to view real-time performance data. Data is only collected while the Top Talkers window is open; it is not persisted in the database.

You can have multiple Top Talker monitors configured simultaneously. You can monitor up to 10 switches for *Fabric mode Top Talkers* and 32 ports and 10 switches for *F-Port mode Top Talkers*; however, you can only monitor one director or port for each Top Talker you configure.

**Note:** Top Talkers requires Fabric OS v6.2 or later and the Advance Performance Monitoring license on the director.

## Starting a Top Talker monitor

To set up a Top Talker monitor, select a director or an F-port on the IBM Network Advisor Product List or Topology Map view, and then select **Monitor** → **Performance** → **Top Talkers** from the menu bar. Alternately, you can right-click a director or an F-port, and select **Performance** → **Top Talkers** (Figure 9-17). If you select a director, the Top Talker monitor starts in Fabric mode. If you select an F-port, the Top Talker monitor starts in F-port mode.

### Fabric mode Top Talkers:

- ▶ A fabric mode Top Talker and an end-to-end monitor cannot be configured on the same fabric. You must delete the end-to-end monitor before you configure the fabric mode Top Talker.
- ▶ A fabric mode Top Talker and an F-port mode Top Talker cannot be configured on the same fabric. You must delete the F-port mode Top Talker before you configure the fabric mode Top Talker.

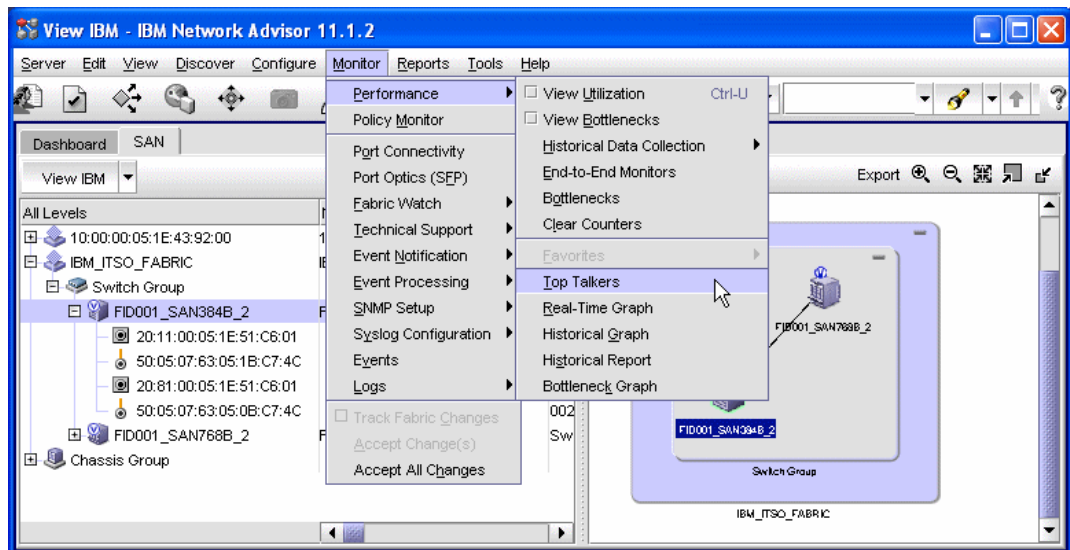


Figure 9-17 Top Talker Monitor launch

The Top Talkers window opens showing, by default, five conversations with the highest Tx+Rx MBps data rate (Figure 9-18 on page 345). The Top Talker monitor is now set up and can be used. You will get the real-time performance data of the selected objects.



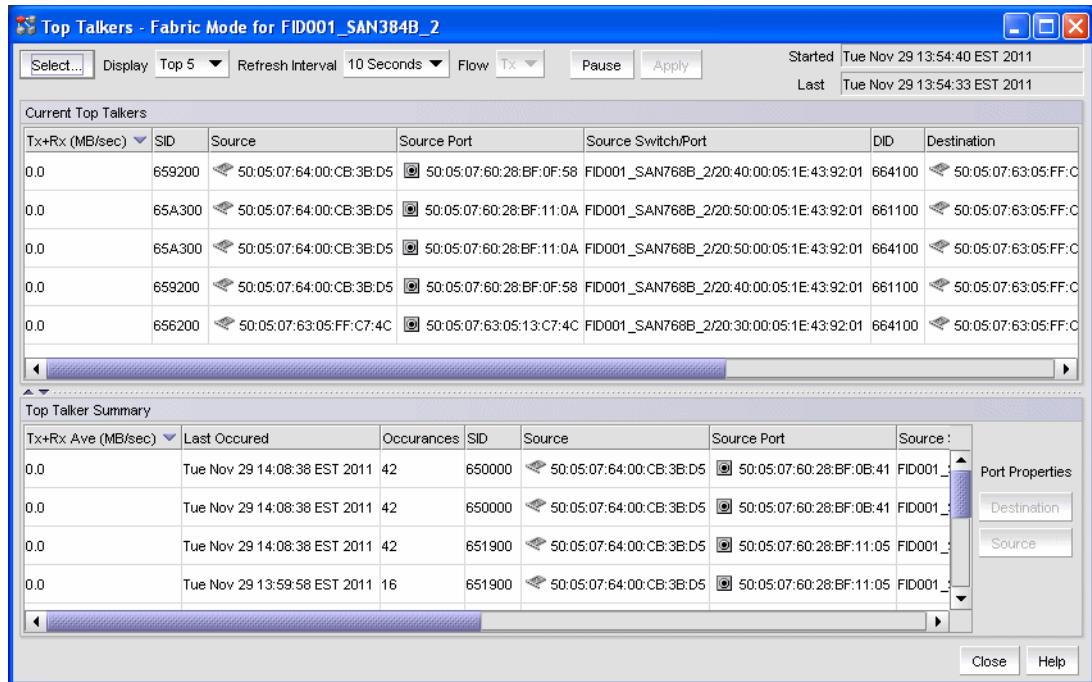


Figure 9-18 Top Talkers window

The Current Top Talkers table (upper part of the window) shows the current most active conversations. The top 20 conversations can be displayed.

The Top Talker Summary table (lower part of the window) shows ports displays all Top Talkers that occurred since the Top Talkers monitor was opened (displays a maximum of 360). When the maximum is reached, the oldest Top Talker drops as a new one occurs.

### Customizing a Top Talker monitor

The Top Talkers window contains buttons that you can use to filter and modify the data displayed on the window:

- ▶ Use the Select button to change the director or port that is monitored. You can also change the Top Talker mode (Fabric mode or F-port mode).
- ▶ Use the Display drop-down menu to select how many entries are displayed in the Current Top Talkers table. Up to 20 conversations can be displayed.
- ▶ Use the Refresh Interval drop-down to set the data refresh interval between 10 seconds and one minute.
- ▶ Use the Flow drop-down to select which metric is monitored. The options are Rx (receive), Tx (transmit), or Tx+Rx MBps. This option is available only for the F-port mode.
- ▶ Click **Apply** to activate the changes in Display, Refresh Interval, and Flow options.
- ▶ Click **Pause** to pause the Top Talker monitor.
- ▶ Click **Close** to delete the Top Talker monitor.

### Exporting and printing Top Talker data

To export Top Talker data to a file for further analysis, right-click anywhere on the Current Top Talker or Top Talker Summary table, and select **Export Table**. This action exports all data in the respective table in a file you are prompted to specify. You also have the option to export

data for a single table row only by selecting **Export Row** or to print the table by selecting **Print**.

## 9.7 Web Tools Performance Monitor

The Web Tools application for IBM System Storage b-type switches includes a performance monitor that you can use to analyze director performance. To access the performance monitor:

1. Launch Web Tools by selecting a director on the IBM Network Advisor Product List or Topology Map view. Select **Configure** → **Element Manager** → **Hardware** from the menu bar. Alternately, you can right-click a director, and select **Element Manager** → **Hardware**.
2. Select **Monitor** → **Performance Monitor** from the Web Tools menu bar, or click **Performance Monitor** in the Tasks list in the left part of the window. The Performance Monitor window opens (Figure 9-19).

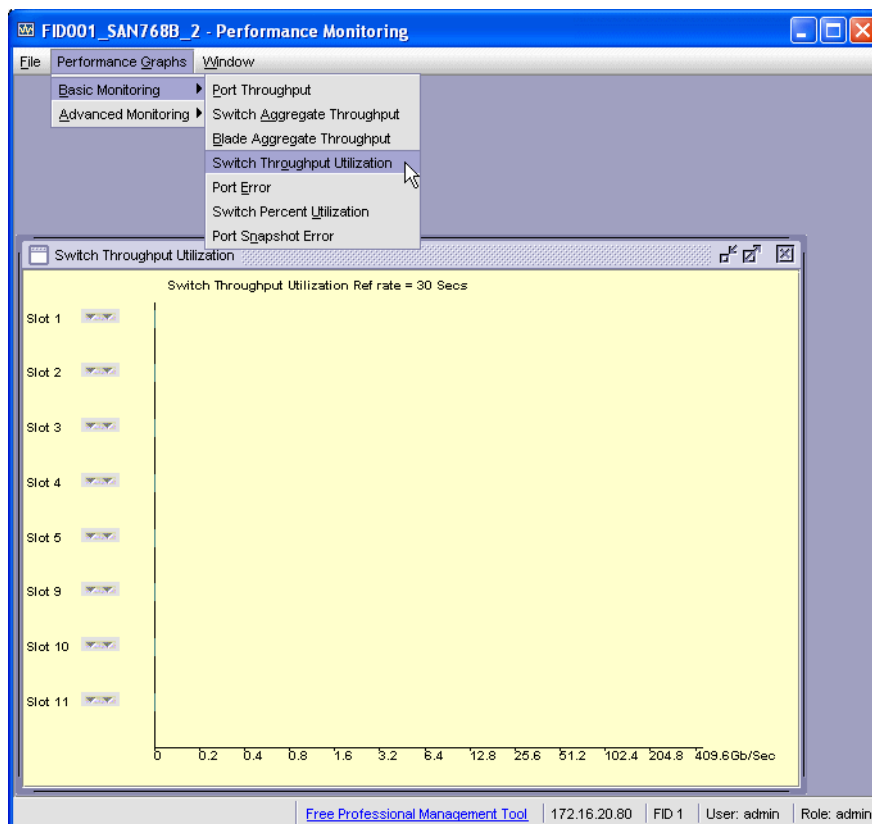


Figure 9-19 Web Tools Performance Monitor

The performance monitor provides the following graphs, which you can view by selecting **Performance Graphs** from the menu bar.

Basic Monitoring graphs:

- ▶ Port Throughput
- ▶ Switch Aggregate Throughput
- ▶ Blade Aggregate Throughput
- ▶ Switch Throughput Utilization
- ▶ Port Error

- ▶ Switch Percent Utilization
- ▶ Port Snapshot Error

Advanced Monitoring graphs:

- ▶ SID/DID Performance
- ▶ SCSI Commands:
  - SCSI Read/Write on a Lun per port
  - SCSI Read on a Lun per port
  - SCSI Write on a Lun per port
  - SCSI Read/Write per port
  - SCSI Read per port
  - SCSI Write per port
- ▶ SCSI versus IP Traffic
- ▶ ALPA Error





## Debugging FICON problems

We describe Problem Determination (PD) tools and techniques that you can use in a FICON environment, which include the following topics:

- ▶ Preparing for problem determination activities:
  - Using the D M - Display Matrix command
  - Creating a CONFIG member
  - IBM Health Checker: Command Response (CMR) Time Function
  - Controlling planned and unplanned changes
  - Tracking some I/O messages
- ▶ Problem determination approach for FICON
- ▶ Diagnostic procedure guideline
- ▶ Environmental record, editing, and printing program
- ▶ FICON link incident reporting
- ▶ FICON Purge Path Extended
- ▶ Helpful z/OS commands
- ▶ Node descriptor
- ▶ IBM Network Advisor PD Tools:
  - Logs
  - SupportSave
  - Port Error Statistics
  - Fibre Channel (FC) Troubleshooting: Trace Route, Device connectivity, and Diagnostic Port test.
  - Helpful Command Line Interface commands
- ▶ HMC PD Tools
- ▶ Common z/OS FICON error message: Interpretation

## 10.1 Preparing for problem determination activities

The problem determination (PD) information provided here assumes that the environment adheres to the architecture specification and implementation levels for the System z servers, the FICON Director, and control units.

We can simplify some problem determination tasks in a z/OS zSeries environment by implementing the procedures written in this section.

### 10.1.1 Using the D M - Display Matrix command

After your system has IPLed, execute a D M command to create an entry in the z/OS system log that will function as a “footprint” to record the processor and I/O hardware configuration at IPL time. We receive status information about Processor, Channel paths, Devices, and Storage.

This command can be automatically performed during the IPL process. The command can be added to member IEACMDxx in SYS1.PARMLIB to automate this process.

The output of this command can be used later during problem determination to help you understand why the status of a hardware resource is in a certain condition.

### 10.1.2 Creating a CONFIG member

A config-ID member, CONFIGxx, can be created in SYS1.PARMLIB using Hardware Configuration Definition (HCD) from an existing I/O Definition (IODF) dataset. When a config-ID member is created in the z/OS PARMLIB, the command D M=CONFIG(xx) is used to show any deviation from the required operating configuration.

The D M=CONFIG(xx) command can be invoked by Operations on a periodic basis (one time every eight hours, at a minimum), and at the operator shift changeover to determine if there are any deviations from the initial configuration status.

The D M=CONFIG(xx) command is also useful after any of the following events:

- ▶ I/O-related repair activity
- ▶ I/O-related physical change activity:
  - Changes to cabling
  - FICON Director settings
  - Add, remove, customize, or recustomize a Control Unit (CU) or device
- ▶ Dynamic I/O re-configuration activity for either software or hardware changes

During scheduled I/O-related change activity, more than one config-ID might be needed to represent the I/O configuration at a particular time or interval.

The D M=CONFIG(xx) command calls member CONFIGxx in SYS1.PARMLIB. The CONFIGxx member contains a predefined hardware configuration that should reflect all the resources critical to system operation. In response to the D M=CONFIG(xx) command, the system compares the contents of the CONFIGxx member and the current configuration. The command displays the possible differences to the operator, who can take the corrective actions, if needed.

Example 10-1 on page 351 shows a CONFIGxx data set member.

*Example 10-1 SYS1.PARMLIB(CONFIGxx)*

---

```
CPU(0),ONLINE
CPU(1),ONLINE
CPU(2),ONLINE
CHP(00-3F),ONLINE
STOR(E=0),ONLINE
STOR(E=1),ONLINE
STOR(E=2),ONLINE
DEV(0200-02FF),ONLINE
DEV(1000-2FFF),ONLINE
DEV(5100-6FFF),ONLINE
DEV(9A00-BF00),ONLINE
...
```

---

Example 10-2 shows the output of a `D M=CONFIG(xx)` command where deviations from a desired configuration (S1) exist. If the current running configuration matches the definitions in the CONFIGxx SYS1.PARMLIB member, the output of the z/OS command contains the message `NO DEVIATION FROM REQUESTED CONFIGURATION`.

*Example 10-2 D M=CONFIG(xx) output*

---

```
D M=CONFIG(S1)
IEE097I 11.38.13 DEVIATION STATUS 568
FROM CONFIGS1
DEVICE   DESIRED   ACTUAL
D2A6     ONLINE    OFFLINE
D080     BASE      HIPERPAV
E200     OFFLINE   ONLINE
F890     PAV      UNBOUND
```

---

### 10.1.3 IBM Health Checker: Command Response (CMR) Time Function

APAR OA33367 (z/OS V1R10 or later) introduces a new IBM Health Check function for IOS. This new function will help us to recognize problems, such as hardware errors, I/O mis-configurations, and congestion by detecting any inconsistent Command Response (CMR) time.

The following messages alert us:

- ▶ IOSHC112E: Appears when function `IOS_CMRTIME_MONITOR` detects a CMR Time exception, see “IOSHC112E” on page 393.
- ▶ IOSHC113I: Message extension of IOSHC112E. Gives us detailed information about the CMR Time exception detected, see “IOSHC113I” on page 393.

For more information about how to configure this new function (`IOS_CMRTIME_MONITOR`), refer to *IBM Health Checker for z/OS User's Guide*, SA22-7994.

### 10.1.4 Controlling planned and unplanned changes

Most common problems are due to configuration changes. Therefore, it is important to be aware of all planned and unplanned changes. Examples of situations that caused problems are listed in this section.

## Problems since last Power-on Reset (POR), Activate, or IPL

The problems since the last POR, Activate, or IPL can be:

- ▶ Wrong Hardware Management Console (HMC) Central Processor Complex (CPC) reset profile used for the CPC for the Activate. This can cause the wrong I/O Configuration Data Set (IOCDS) to be selected and used.
- ▶ Wrong IOCDS associated to the Reset Profile used to activate.
- ▶ I/O configuration definition changes performed that do not match the physical configuration.
- ▶ Wrong LOAD profile used, or wrong LOAD device or IODF device specified in the LOAD profile:
  - The wrong LOAD device used for the IPL can cause the wrong OSCONFIG to be used. (The OSCONFIG member contains the z/OS software I/O configuration definition.)
  - The wrong IODF device (volume) used can point to a different I/O configuration.

**Note:** For information related to IODF configuration, such as OSCONFIG member, refer to 5.5, “Building and activating the IODF” on page 137.

## Problems since last physical configuration change

The problems since the last physical configuration change can include:

- ▶ Physical cabling changes at System z Server, FICON Director, patch panels, DWDMs or Control Unit side.
- ▶ Addition or removal of FICON Directors or control units.

## Problems since last control unit customization changes

Problems since the last FICON Director changes include:

- Matrix changes (prohibit/allow)
- Customization changes: Fabric, FICON Director, or Port bindings
- Port status changed: ports being blocked/disable or taken offline
- Zoning customization changes
- Port cabling changes

## Problems since last software/hardware code level changes

The problems that can occur since the last software/hardware code level changes are:

- ▶ z/OS Maintenance (PTFs)
- ▶ Central Electronic Complex code level changes (Driver, Bundle, MCLs)
- ▶ FICON Director firmware (FOS)
- ▶ Control Unit (CU) firmware

### 10.1.5 Tracking some I/O messages

Another good practice is to track some I/O Messages. Keeping track of these messages allows us to detect errors that can potentially be a hardware fault, such as a fiber problem.

Examples of I/O Messages to track are:

- ▶ IOS050I, IOS054I: Reporting Channel Detected Errors
- ▶ IOSHC112E, IOSHC113I: Reporting inconsistent Command Response (CMR) time from one or more control unit



For more information about these messages, see 10.11, “Common z/OS FICON error message: Interpretation” on page 377.

More than one error message can appear over time for the same path without any known reason. Depending on the amount of time between messages, we must perform some diagnostics on the affected path to validate it.

Throughout this chapter we will see where we can look for information that can help us on finding details of the affected path.

## 10.2 Problem determination approach for FICON

Before you initiate any problem determination (PD) tasks, you must have a clear understanding of the problem or condition. Begin by taking note of IOS messages that are issued (see 10.11, “Common z/OS FICON error message: Interpretation” on page 377 for examples). IOS messages provide useful information and are an excellent starting point for problem determination. Determining the status of all components in the affected I/O path is also key. Techniques to use for component status determination are shown in 10.9, “IBM Network Advisor PD Tools” on page 364, Appendix E, “Using HMC and SE for problem determination information” on page 425, and in Appendix F, “Useful z/OS commands” on page 471.

You also need current and accurate information and documentation to resolve problem situations, so try to have the information listed here available:

- ▶ The symptom or z/OS message associated with the problem that needs to be addressed.
- ▶ The device number or device numbers of the problem device or devices.
- ▶ The CHPID numbers of the problem channels or the channel paths to the device or devices.
- ▶ The CPC and logical partition names that the z/OS is running on.
- ▶ A configuration I/O connectivity complete diagram (all paths and FICON Director connections for a multi-path device) of the affected device.
- ▶ The physical location of all involved hardware equipment and their respective consoles, when applicable.
- ▶ For a FICON CTC-related problem:
  - Target CPC
  - Target Channel Subsystem
  - Target Image (LPAR)
  - Target matching device number (or the source/target CTC device pair)

z/OS offers an extensive list of commands that will provide most of the required PD information listed. The HMC and Support Element (SE) panels include specific Problem Determination functions that might also be required to use when investigating FICON channel problems. IBM Network Advisor also has different tools for debugging. z/OS commands, HMC/SE problem determination facilities, and IBM Network Advisor PD Tools are discussed in Appendix F, “Useful z/OS commands” on page 471 and Appendix E, “Using HMC and SE for problem determination information” on page 425, and 10.9, “IBM Network Advisor PD Tools” on page 364.

For FICON channel problem determination, a significant amount of information is available in various areas. The following sections provide an overview of what can be used based on four

levels of tools. Figure 10-1 illustrates the levels of software and hardware and the tools provided from a verification perspective.

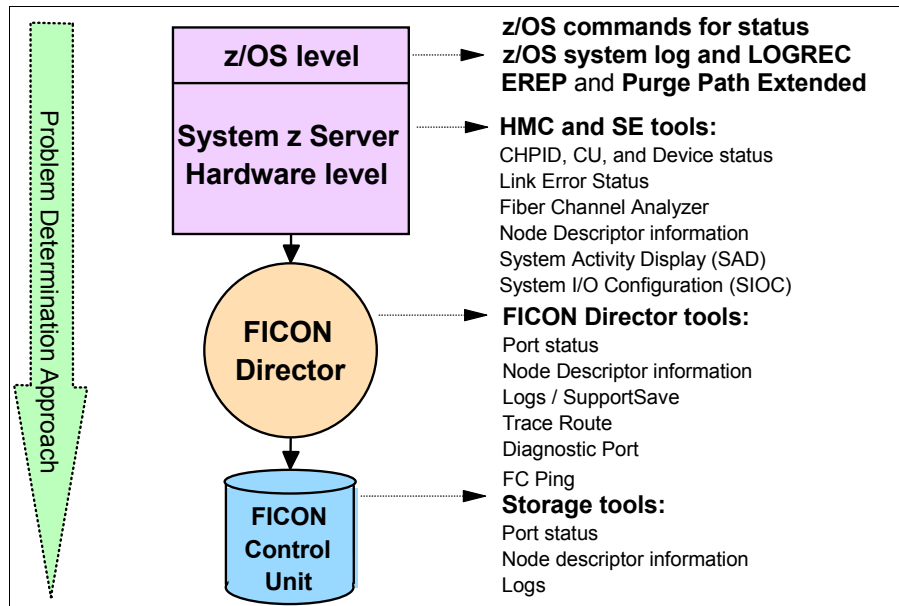


Figure 10-1 Sources of information

## Most common FICON failures

There are two common failures seen on FICON channels:

- ▶ Failure to complete initialization or establish a logical path.
- ▶ Failure detected in the communications of command and data by either the channel or the control unit.

Initialization failures might or might not create a Hardware Message on the HMC/SE. The SE Problem Determination panels, shown under “Displaying the Analyze Serial Link Status PD panel” on page 447, will clearly show initialization failures on a per-FICON channel basis. There is also a panel available under the Service logon user ID called “IFCC and Other Errors” that will display information related to initialization types of failures. Refer to E.4, “Displaying FICON IFCC and Other errors” on page 453.

Channel Control Check (CCC) errors, initialization failures, Channel Detected Errors and Interface Control Checks (IFCC) can be caused by failures in FICON hardware or in the FICON definition setup.

## FICON hardware-related failures

The FICON hardware-related failures are:

- ▶ FICON channel card can cause initialization IFCC and CCC failures.
- ▶ FICON I/O Control Unit (CU) Host Bay Adapter (HBA) can cause initialization and IFCC failures.
- ▶ FICON Director can cause initialization and IFCC failures.
- ▶ Fibers in general can cause initialization and IFCC failures.

## FICON definition-related failures

The FICON definition-related failures are:

- ▶ I/O Configuration Program (IOCP) definition errors can cause initialization failures.
- ▶ Fabric definitions can cause initialization and IFCC failures.
- ▶ Control Unit definition errors can cause initialization and IFCC failures.

## FICON Zoning/Traffic Isolation (TI) Zoning-related failures

Zoning issues in fabrics can cause problems that are difficult to diagnose. If zoning or TI zoning was implemented, the zoning tables in the fabric must be updated whenever a new port is installed and connected to a zoned fabric.

In case of having implemented the Allow/Prohibit Matrix, you might have to review its configuration in case of adding new ports.

# 10.3 Diagnostic procedure guideline

An example of a diagnostic procedure is outlined here. It is difficult to handle all problems that can arise because of the multiple factors that can occur.

Use this procedure as a general guidance.

**Note:** During this procedure, there are different internal and external references to help you find information related.

### Problem awareness

*Something seems to happen: first information I receive...:*

- ▶ Users complain
- ▶ Operator advises of red messages
- ▶ Performance analyst notices something unusual
- ▶ ...

### Symptom gathering

*What is happening? Let's look through the data we have...:*

- ▶ Message Information: JOB Name, Channel command, CHPid, Device,...
- ▶ ABEND information: Possible DUMP generated. Refer to "z/OS MVS System Completion Codes" (SA22-7626) and "z/OS MVS Dump Output Messages" (SA22-7590).
- ▶ zOS system/loop/jobs hanging: Use the MVS dump command or, if no recovery is possible, use the Stand-Alone Dump (SADMP). Refer to "z/OS MVS System commands" (SA22-7627) and "z/OS MVS Diagnosis: Tools and Service Aids" (GA22-7589).
- ▶ Other information: CPC and Logical Partition (LPAR) names where the z/OS is running on, physical location of involved hardware equipment.

**Note:** It is important to prepare a connectivity map of the involved components as soon as you have them identified. It helps you focus the problem during the PD tasks.

## Collect diagnostic data

*What other information can help me?...:*

- ▶ System LOG (SYSLOG): Look for other related messages. Refer to 10.11, “Common z/OS FICON error message: Interpretation” on page 377.
- ▶ EREP: Print LOGREC records to find related information. Refer to 10.4, “Environmental record, editing, and printing program” on page 357.
- ▶ HMC: Look for related Hardware messages, IFCC or link errors. Refer to E.4, “Displaying FICON IFCC and Other errors” on page 453.
- ▶ IBM Network Advisor (for FICON Director): Look for related messages. Refer to 10.9.1, “Logs” on page 364.
- ▶ Check Hardware configuration (e.g. through HCM). Refer to E.6, “Displaying the System I/O Configuration Analyzer panel” on page 469.
- ▶ DUMP: Use MVS DUMP command, or set a SLIP, or take a SADMP (SADMP in case the system is unrecoverable). Refer to “z/OS MVS System commands” (SA22-7627) and “z/OS MVS Diagnosis: Tools and Service Aids” (GA22-7589).
- ▶ Traces: Use GTF Trace if the need is determined after some analysis. Refer to “z/OS MVS System commands” (SA22-7627) and “z/OS MVS Diagnosis: Tools and Service Aids” (GA22-7589).

## Analyze diagnostic data and symptoms

*Let's do “Hands on lab”...:*

- ▶ Determine the status of all components in the affected I/O path, if already identified:
  - Display of Devices, CHPIDs, paths and switches; if already identified. Refer to 10.7, “Helpful z/OS commands” on page 360.
  - Use HMC PD Tools to check the status of the channel, fiber. Refer to E.3, “Displaying individual FICON channel information” on page 435 and E.5, “Displaying the Fibre Channel Analyzer panel” on page 466.
  - Use IBM Network Advisor (for FICON Director) to review port error statistics and perform diagnostics if necessary (e.g. fcping, trace route, D\_port). Refer to 10.9.3, “Port Error Statistics” on page 368 and 10.9.4, “Fibre Channel (FC) Troubleshooting” on page 370.
- ▶ DUMP: Get the device involved and path.
- ▶ EREP: Look for possible software errors correlated to device errors (by time stamp).
- ▶ GTF Trace: Check if GTF Trace record “Program-controlled input/output interruption (PCI)” shows a FICON Device. Refer to “z/OS MVS Diagnosis: Tools and Service Aids” (GA22-7589).
- ▶ Identify the error type:
  - Is I/O related error? If no, do “normal” PD Processor Support Interface (PSI).
  - Is using a FICON path? If no, do “normal” PD Processor Support Interface (PSI).
  - If you are still not sure, return to step: *Collect diagnostic data*.
- ▶ Messages: Review the actions indicated on the related messages.
- ▶ Abend: Continue with “normal” PD/PSI to determine FICON involvement.
- ▶ Performance problem: run GTF and look for PCI trace records or check sources code.

## Cause identification

*What is my “Root Cause”?...:*

- ▶ No cause identified: Think on likely scenarios and return to *Collect diagnostic data*.
- ▶ Cause identified:
  - FICON related: Continue with *Fix identification and testing*.
  - Non-FICON cause: Continue with “normal” PD/PSI.

## Fix identification and testing

*Do I have FIX?...:*

- ▶ Yes, we have a FIX: Check if we can test. Yes or no, contact the component owner.
- ▶ No, we do not have a FIX: Contact the component owner.

## 10.4 Environmental record, editing, and printing program

When an error occurs, the system records information about the error in the LOGREC data set or the LOGREC log stream. This information provides you with a history of all hardware failures, selected software errors, and selected system conditions. Use the Environmental Record, Editing, and Printing program (EREP) to print system records, to determine the history of the system or to learn about a particular error.

For further information about this topic, refer to *z/OS MVS Diagnosis: Tools and Service Aids*, GA22-7589.

## 10.5 FICON link incident reporting

To improve the ability to capture data for link error analysis, z/OS can register to receive FICON link incident records, which are referred to as Link Maintenance Information (LMI) records. z/OS will display this information about consoles and save it in the system log and LOGREC:

- ▶ IOS580E/IOS581E: System Log message reporting a link incident error. For additional information refer to *MVS System Messages*, SA22-7639.
- ▶ Link Maintenance Information Detail Edit (A2) Report: LOGREC record. For additional information refer to *Environmental Record Editing and Printing Program (EREP): Reference*, GC35-0152.

A link incident record has detailed node information that is gathered for particular link-degraded and link-failure incidents. The node descriptor information is gathered during the Fabric Login or Port Login when the logical path is established.

For additional information pertaining to LMI records, refer to *z/OS MVS Diagnosis: Reference*, GA22-7588.

You can also find information about Link Incident errors on HMC/SE FICON Link Incident log. See E.5, “Displaying the Fibre Channel Analyzer panel” on page 466.

## 10.6 FICON Purge Path Extended

When the extended subchannel logout facility is installed and enabled, all Fibre Channel (FC) channel paths can create extended subchannel logouts, under applicable error conditions,

and generate I/O interruptions that indicate, for example, a logout is pending retrieval. If a logout is pending for an I/O interruption, IOS will retrieve the logout data and include the data in the subchannel logout record that it records in SYS1.LOGREC.

**Note:** FICON Purge Path Extended (PPE) feature was introduced with z10 as a new Availability/Reliability benefit. FICON Purge Path Extended must be supported by the control unit and FICON Director to receive error counters. Check with your product vendor for details.

With FICON PPE enabled, Link Error Statistical Buffer (LESB) information is collected from each Fibre Channel port in the path from host to control unit. This information then results in the abort reason code being built. The channel subsystem passes this information to z/OS as Extended Subchannel Logout Data (ESLD), and an SLH record is cut in the LOGREC. The overall process is shown in Figure 10-2.

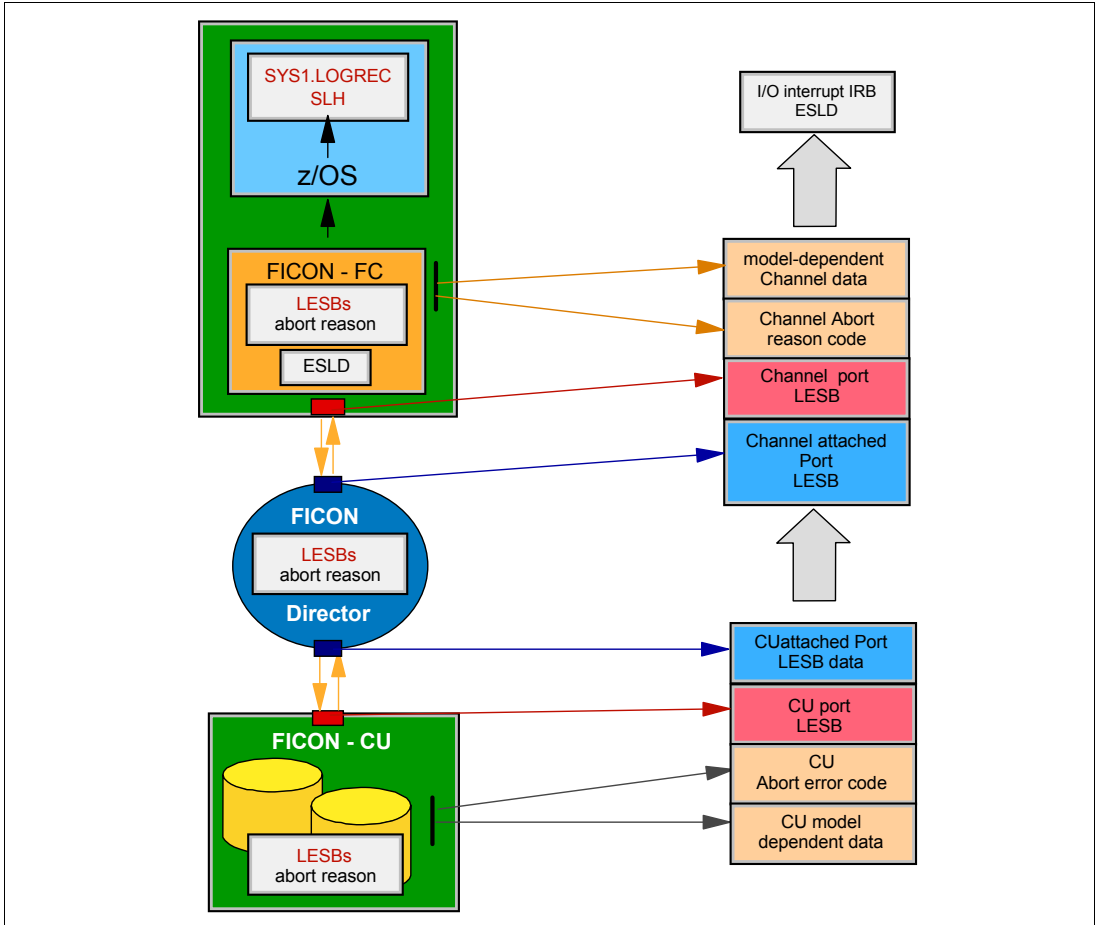


Figure 10-2 FICON Purge Path Extended error recording

This function provides advanced problem determination of Fibre Channel bit errors. A subset of a sample Subchannel Logout Handler (SLH) record, formatted by EREP, is included in Figure 10-3 on page 359.

```

DEVICE NUMBER: 07025 REPORT: SLH EDIT DAY YEAR JOB
                SCP: VS 2 REL. 3 DATE: 299 11
DEVICE NED: 002107.900.IBM.75.0000000L3001.0000
DEVICE TYPE: 3390 CPU MODEL: 2817XA HH MM SS.TH
CHANNEL PATH ID: 45 LOGICAL CPU ID: D52817 TIME: 12 21 29.81
PHYSICAL CHAN ID: 0371 PHYSICAL CPU ADDRESS: 0E
                CC CA FL CT
FAILING CCW E7 7C997AA4 60 0044 VOLUME SERIAL
                SUBCHANNEL ID NUMBER
                ERROR TYPE
SCSW K FLAGS CA US SS CT
04 C24017 7C997AF0 00 02 0000

---UNIT STATUS--- SUB-CHANNEL STATUS -----SCSW FLAGS-----
                FLAG 0 FLAG 1 FLAG 2
ATTENTION 0 PGM-CTLD IRPT 0 CCW FORMAT 1 RESERVED 0 SUBCHANNEL ACTIV 0
STATUS MODIFIER 0 INCORRECT LENGTH 0 PRE-FETCH CCW 1 SSCH FUNCTION 1 DEVICE ACTIVE 0
CONTROL UNIT END 0 PROGRAM CHECK 0 INIT STATUS 0 HSCH FUNCTION 0 SUSPENDED 0
BUSY 0 PROTECTION CHECK 0 ADDR LIMIT 0 CSCH FUNCTION 0 ALERT STATUS 1
CHANNEL END 0 CHAN DATA CHECK 0 SUPP SUSPEND INT 0 RESUME PENDING 0 INTERMED STATUS 0
DEVICE END 0 CHAN CTL CHECK 0 ZERO COND CODE 0 START PENDING 0 PRIMARY STATUS 1
UNIT CHECK 0 I/F CTL CHECK 1 EXTENDED CONTROL 1 HALT PENDING 0 SECONDARY STATUS 1
UNIT EXCEPTION 0 CHAINING CHECK 0 PATH NOT OPER 0 CLEAR PENDING 0 STATUS PENDING 1

----SOFTWARE RECOVERY STATUS-----
HARD FAIL 0
DEGRADE FAIL 0
SOFT FAIL 1
PASSED 0

CHANNEL ERROR ANALYSIS
IRB STORED BY INTERRUPT
TERMINATION BY -- SELECTIVE RESET -- CODE 2
SEQ CODE *** INVALID ***
VALIDITY OF RECORDED DATA
COUNT INVALID
TERMINATION CODE VALID
SEQUENCE CODE INVALID
DEVICE STATUS INVALID
CCW ADDRESS VALID
DEVICE NUMBER VALID
SENSE DATA NOT STORED

EXTENDED SUBCHANNEL LOGOUT DATA
CHANNEL LOGOUT DATA
N-PORT LINK ERROR STATUS BLOCK
LINK FAILURE COUNT: 00000002 LOSS OF SYNCHRONIZATION COUNT: 00000001
LOSS OF SIGNAL COUNT: 00000000 PRIMITIVE SEG PROTOCOL ERROR: 00000000
INVALID TRANSMISSION WORD: 00001295 INVALID CRC COUNT: 00000000

FABRIC ENTRY PORT LINK ERROR STATUS
LINK FAILURE COUNT: 00000000 LOSS OF SYNCHRONIZATION COUNT: 00000000
LOSS OF SIGNAL COUNT: 00000000 PRIMITIVE SEG PROTOCOL ERROR: 00000000
INVALID TRANSMISSION WORD: 00024CCD INVALID CRC COUNT: 00000000
ERROR CODE: 0C - Receive ABTS

MODEL DEPENDENT DATA:
0000 0C000000 9000A800 A8000000 00000000 00000000 00979F00 0000872B 10000005
0020 1E363CBA 10000005 1E363CBA 50050760 1F4058F9 50050764 00C29B46 0006000C
0040 00000000 00000000 08000042 10400000 22800022 00000000 02000000 18100020
0060 50050763 041885C1 50050763 04FFC5C1 00310100 30303231 30373933 3249424D
0080 37353030 30303030 30434347 48310302 00000000 00000000 00000000 00000000
00A0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00C0 00000000 00000000 00000000 00000000

CONTROL UNIT LOGOUT DATA
N-PORT LINK ERROR STATUS BLOCK
LINK FAILURE COUNT: 00000000 LOSS OF SYNCHRONIZATION COUNT: 00000000
LOSS OF SIGNAL COUNT: 00000000 PRIMITIVE SEG PROTOCOL ERROR: 00000000
INVALID TRANSMISSION WORD: 00000000 INVALID CRC COUNT: 00000000

FABRIC ENTRY PORT LINK ERROR STATUS
LINK FAILURE COUNT: 00000000 LOSS OF SYNCHRONIZATION COUNT: 00000000
LOSS OF SIGNAL COUNT: 00000000 PRIMITIVE SEG PROTOCOL ERROR: 00000000
INVALID TRANSMISSION WORD: 00000000 INVALID CRC COUNT: 00000000
ERROR CODE: 00 - Error code transfer not supported

MODEL DEPENDENT DATA:
0000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0020 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0040 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0060 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
0080 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00A0 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00000000
00C0 00000000 00000000 00000000 00000000

HEX DUMP OF RECORD
HEADER 23831800 00030000 0011299F 12212981 00139B46 28178000
0018 D7C4C2F2 C4C2D4F1 E7600044 7C997AA4 801F2032 00000000 40046780 04C24017
...

```

Figure 10-3 Formatted Subchannel Logout Handler record

Figure 10-4 shows the CHANNEL and CONTROL UNIT LOGOUT DATA fields extracted from the EXTENDED SUBCHANNEL LOGOUT DATA on EREP SLH record from Figure 10-3 on page 359. The pictures shows you what information you can look for: Link Error Status Block (LESB) counters and PPE error code.

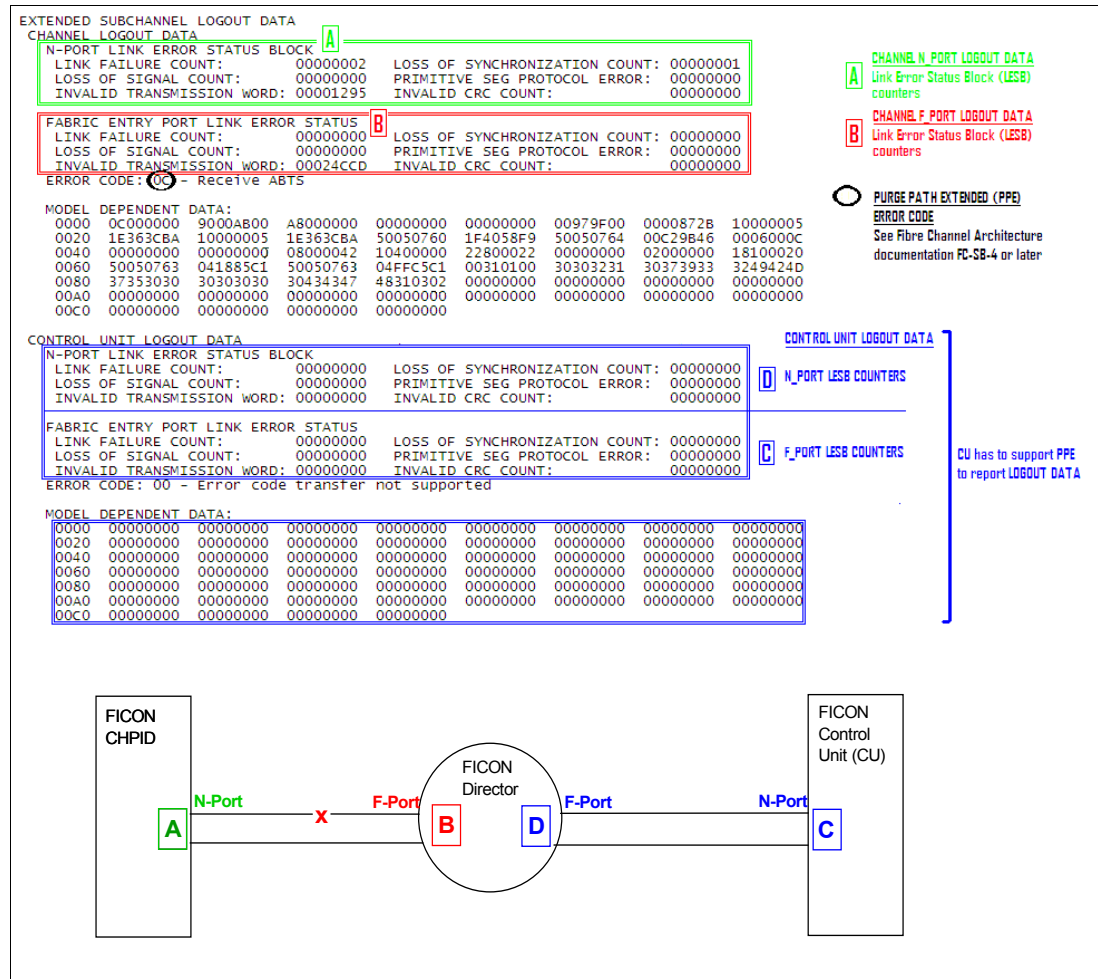


Figure 10-4 CHANNEL and CONTROL UNIT LOGOUT details

**Note:** Introduced in z/OS V1R13, SLH EREP record format improves LESB error counter and PPE error code information display.

For a detailed description regarding the debugging of these records, refer to the ANSI standard architecture, which can be found at:

<http://www.t11.org>

HMC/SE also gives information about Link Error Statistics Block (LESB). See Appendix E, “Displaying the Analyze Link Error Statistics Block PD panel” on page 450.

## 10.7 Helpful z/OS commands

The display results of the following commands, if issued from the affected logical partition (LPAR), can be helpful for documenting the configuration and analyzing the errors in a FICON



environment. Therefore, before you begin to perform problem determination that involves a FICON Director, use these commands. You also need to obtain a copy of the z/OS syslog.

- ▶ D M=CHP displays the status and type of CHPIDs
- ▶ D M=CHP(chp) chp is the affected CHPID
- ▶ D U,,,devn,1 devn is the affected device address
- ▶ DS P,devn devn is the affected device address
- ▶ D M=DEV(devn) devn is the affected device address
- ▶ D M=DEV(devn,(chp)) devn is the affected device address and chp is the failing CHPID.

**Note:** Introduced in z/OS 1.13, command D M=DEV(devn,(chp)) now shows the reason for an offline path.

- ▶ D M=SWITCH(xx) displays the ports configured and in use by the specified FICON Director number.

Examples of the z/OS commands can be found in Appendix F, “Useful z/OS commands” on page 471.

Example 10-3 shows the output of a z/OS D M=SWITCH(xx) command.

*Example 10-3 Output of the D M=SWITCH(65) z/OS command*

```

D M=SWITCH(65)
IEE174I 12.45.15 DISPLAY M 215
SWITCH 0065, PORT STATUS
      0 1 2 3 4 5 6 7 8 9 A B C D E F
0    u u u c u u u u u u u u u u u u
1    u u u u u u u u u u u c u u u u
2    u u u u u u u u u u u u u x u u
3    u u u u x u u u u u u u u u u u
4    . . . . . . . . . . . . . . .
5    . . . . . . . . . . . . . . .
6    . . . . . . . . . . . . . . .
7    . . . . . . . . . . . . . . .
8    . . . . . . . . . . . . . . .
9    . . . . . . . . . . . . . . .
A    . . . . . . . . . . . . . . .
B    . . . . . . . . . . . . . . .
C    . . . . . . . . . . . . . . .
D    . . . . . . . . . . . . . . .
E    . . . . . . . . . . . . . . .
F    . . . . . . . . . . . . . . .
***** SYMBOL EXPLANATION *****
+ DCM ALLOWED          - DCM NOT ALLOWED BY OPERATOR
x NOT DCM ELIGIBLE    p DCM NOT ALLOWED DUE TO PORT STATE
c CHANNEL ATTACHED    $ UNABLE TO DETERMINE CURRENT ATTACHMENT
u NOT ATTACHED        . DOES NOT EXIST

```

The D M=SWITCH(xx) command provides important information about the configured and in-use ports at the specified switch, along with their current state.

Notice in Example 10-3 that Switch 65, ports 03 and 1B are in use and attached to a FICON Channel. Ports 2D and 34 display an x state. These ports are used as ISL ports.

## 10.8 Node descriptor

The node descriptor will provide you with useful information such as the manufacturer, node type, node status, serial number, Tag information, WWNN, and WWPN.

The following topics explain two ways to obtain Node Descriptor information:

- ▶ 10.8.1, “View Node Descriptors from HMC/SE” on page 362
- ▶ 10.8.2, “View Node Descriptors from IBM Network Advisor” on page 362

**Note:** The specific node-identification data is referred to as the Node Descriptor (ND); 32-byte field that describes a node. For more information refer to FC-SB-4 manual, which can be found at: <http://www.t11.org>

For more information about Node Descriptor (ND), refer to FC-SB-4 manual, which can be found at:

<http://www.t11.org>

### 10.8.1 View Node Descriptors from HMC/SE

For an explanation about how to obtain this information from a System z server, refer to topic “Displaying the Analyze Channel Information panels” on page 440 on Appendix E, “Using HMC and SE for problem determination information” on page 425.

### 10.8.2 View Node Descriptors from IBM Network Advisor

If not already done, to change the IBM Network Advisor view configuration to FICON:

1. Click **SERVER** → **Options**.
2. In the Option window, select **SAN Display**, and click **Setup FICON Display**.
3. Click **OK** to close the window and save the changes.

You will see the Node Descriptor details in the IBM Network Advisor for each port at the left side of the window. To change the port view, right-click a Director (on the left of the window) and select **Port Display**. Select the details you want to see.

To also see the Node Descriptor for a Director:

1. Right-click and select **Element Manager** → **Hardware**.
2. In the Directors Element Manager, click **Name Server** at the left side of the window to open the Name Server window displaying all connected devices, as shown in Figure 10-5 on page 363.

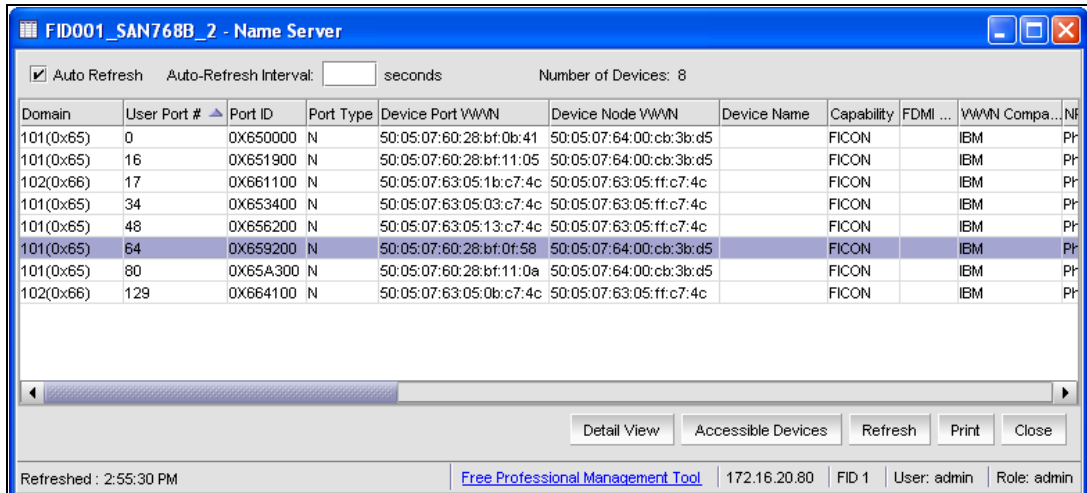


Figure 10-5 Node Descriptor overview

Clicking the **Accessible Devices** button displays all devices that have access to the selected list entry, that is: each device in the same Zone as the selected list entry.

Selecting the list entry and clicking **Detail View** shows us a detailed view for the selected device, as shown in Figure 10-6.

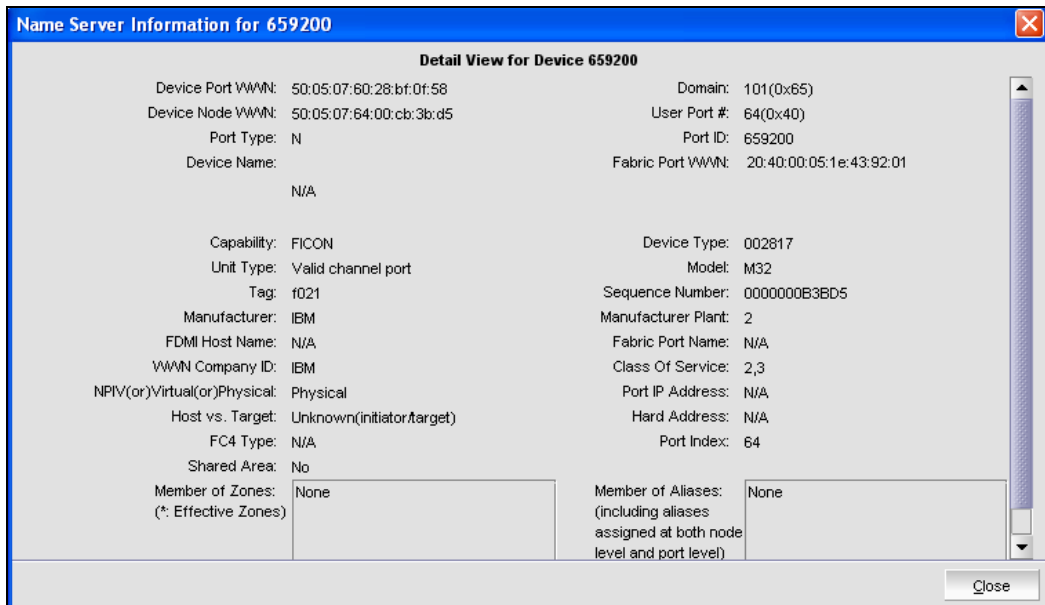


Figure 10-6 Node descriptor detailed view

The Tag field for the System z server is a four-digit number, and it consists of two parts. The last two digits are the CHPID and the first digit gives you information about the Channel Subsystem (CSS) in which the CHPID is defined. In our example, the Tag is f021, which means that the CHPID is 21 and is defined in CSS 0, 1, 2, and 3 (see Table 10-1).

Table 10-1 CSS mapping

| First digit | CSS | First digit | CSS     | First digit | CSS         |
|-------------|-----|-------------|---------|-------------|-------------|
| 1           | 3   | 6           | 1 and 2 | B           | 0, 2, and 3 |

| First digit | CSS     | First digit | CSS         | First digit | CSS            |
|-------------|---------|-------------|-------------|-------------|----------------|
| 2           | 2       | 7           | 1, 2, and 3 | C           | 0 and 1        |
| 3           | 2 and 3 | 8           | 0           | D           | 0, 1, and 3    |
| 4           | 1       | 9           | 0 and 3     | E           | 0, 1, and 2    |
| 5           | 1 and 3 | A           | 0 and 2     | F           | 0, 1, 2, and 3 |

You can also use the Command Line Interface (CLI) of the FICON Director to display Node Descriptor information. Related CLI command: **ficonshow rnid**. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

The Tag field details for a DS8000 are shown in *Configuring DS8000 I/O ports - “I/O port layout”* on page 417.

## 10.9 IBM Network Advisor PD Tools

Throughout this chapter we show the main IBM Network Advisor PD Tools:

- ▶ Logs
- ▶ SupportSave
- ▶ Port Error Statistics
- ▶ Fibre Channel (FC) Troubleshooting:
  - FC Trace Route
  - Device Connectivity
  - Fabric Device Sharing
  - Diagnostic Port Test

### 10.9.1 Logs

The IBM Network Advisor has several logs that include information which will help you during problem determination. The Master Log is displayed in the lower part of the IBM Network Advisor and it includes, by default, all events. You can also set a filter by clicking **Filter** and selecting the events you want to show in the Master Log.

If you want a specific log, select **Monitor** → **Logs**, as shown in Figure 10-7 on page 365.

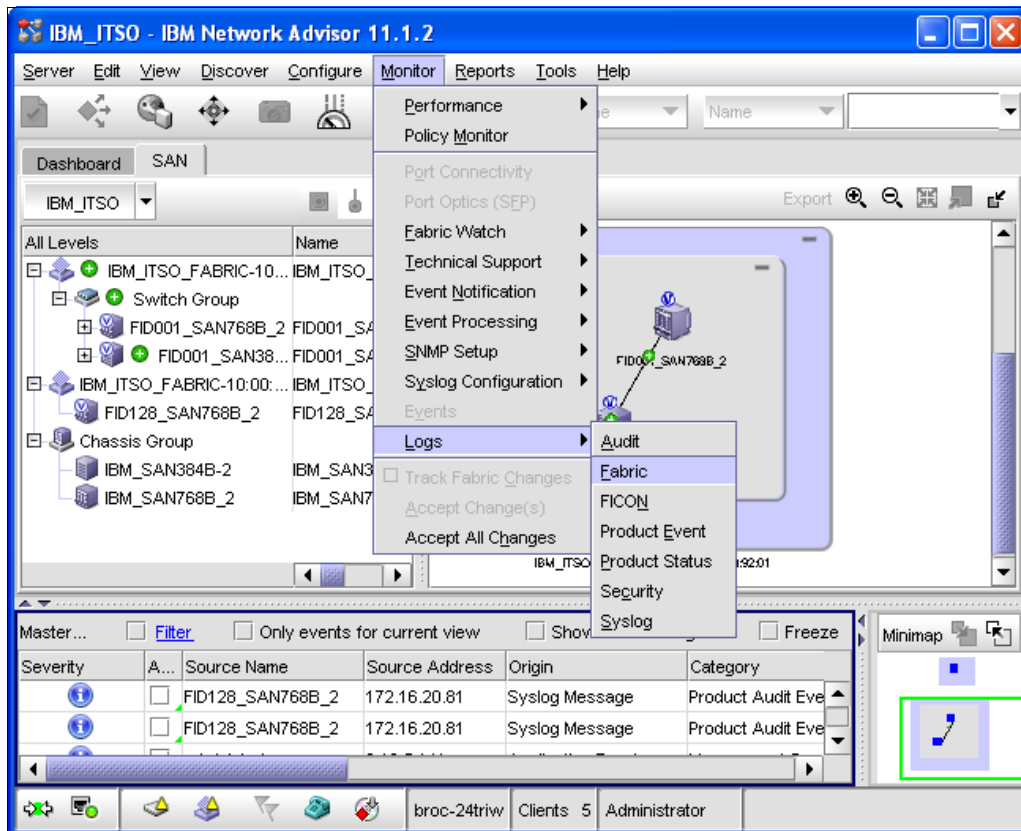


Figure 10-7 Selecting a Log

You can display the log from a specific Director or logical Switch by right-clicking the Director and selecting **Events**.

Related CLI commands: **errdump** and **auditdump**. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

## 10.9.2 SupportSave

To perform a Director or Logical Switch SupportSave, select **Monitor** → **Technical Support** → **Product / Host SupportSave** in the IBM Network Advisor menu bar. The Technical Support Data window appears, as shown in Figure 10-8 on page 366.

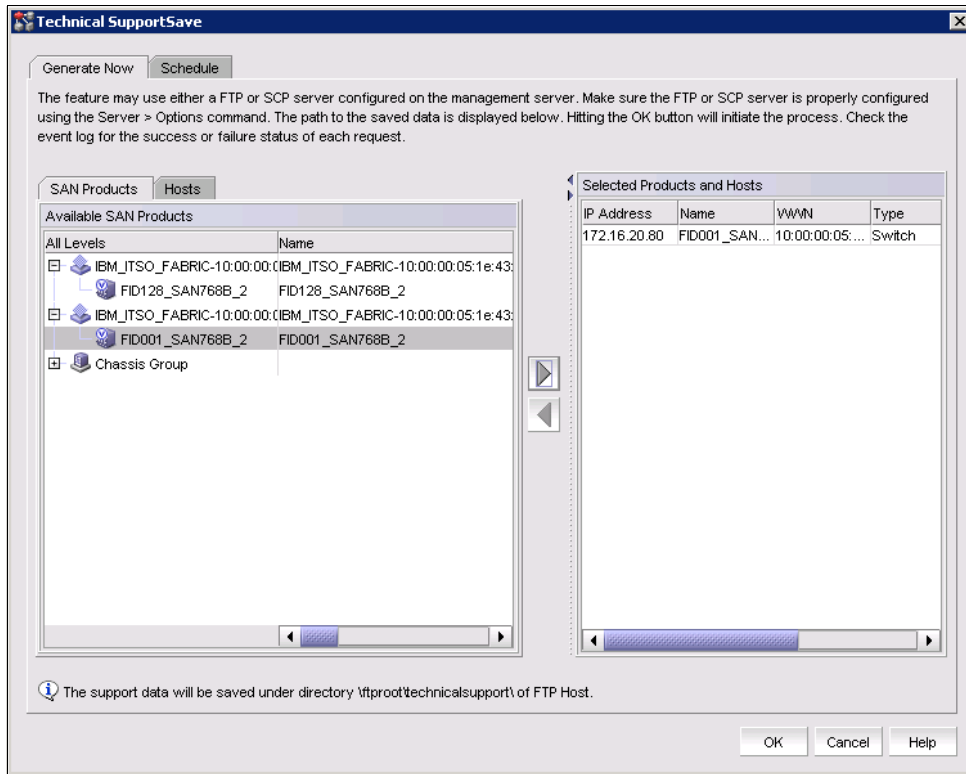


Figure 10-8 Product / Host SupportSave panel

**Note:** Product / Host SupportSave panel gives you the option to perform the SupportSave at this moment, *Generate Now* tab, or to perform the process with a planned frequency, *Schedule* tab.

Select the Director or logical switch at the left side of the window and click the arrow pointing to the right, to put it in the list of selected switches for data capture. Click **OK** to start the data capture process.

The Technical SupportSave Status window will appear showing the progression of the process, as shown on Figure 10-9. The process can take up to 30 minutes, depending on the size of your Director configuration.

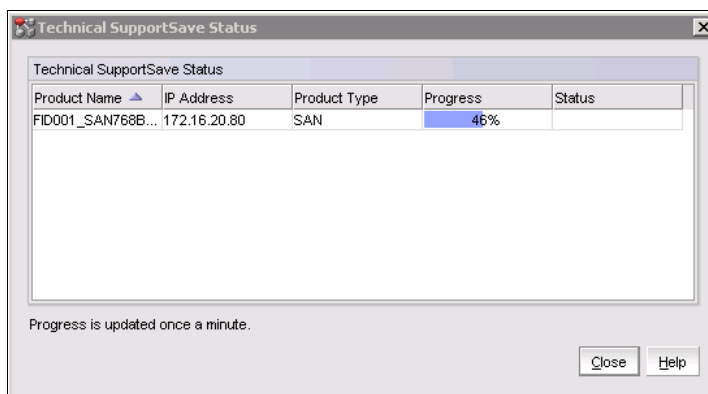


Figure 10-9 Technical SupportSave Status panel

Checking the Master Log messages is another way to validate the successful end of the SupportSave process. See Figure 10-10.




Figure 10-10 Master Log SAN Product SupportSave messages

After SupportSave process finishes, you can select the resultant archive file by going to **Monitor** → **Technical Support** → **View Repository**. The SupportSave file for a Director or a Logical switch will have “SAN Product” value in “Type” column. Select the correct archive file by looking at the date and time stamp included in the file name, as shown in Figure 10-11.

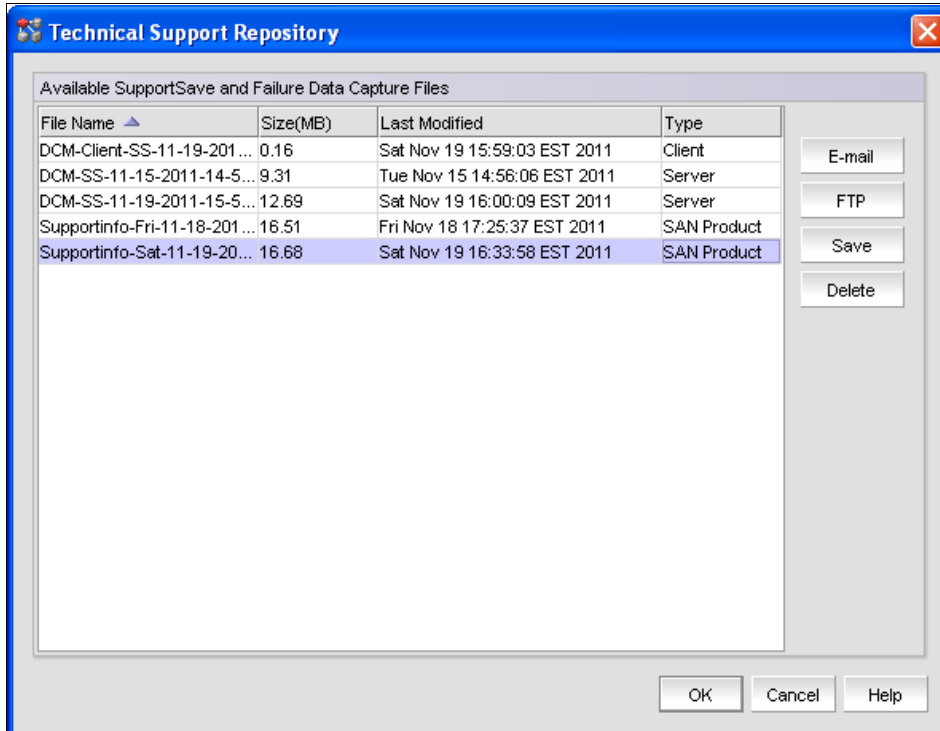


Figure 10-11 Technical Support Repository panel

Technical Support Repository panel gives you the following options:

- ▶ E-mail, to send captured data file to a specific email address.
- ▶ FTP, to send captured data file to an external FTP server.
- ▶ SAVE, to save captured data file in a specific location, for example to your local hard disk, or a USB device.
- ▶ DELETE, to erase useless SupportSave data files. When using the Delete option, you will receive a confirmation panel after you click **OK**. The confirmation panel will give you the possibility of undo the delete actions.

If you are at the IBM Network Advisor Server, you can also go into the following directory to get the Director or Logical Switch archive data file:

C:\Program Files\IBM Network Advisor 11.1.2\data\ftproot\technicalsupport

SupportSave is finished now and can be sent to the Support Center.

**Note:** SupportSave gathers RASLOG, TRACE, SupportShow, core file, FFDC data, and other support information. You can use CLI command SupportShow when, for some reason, you are not able to take the SupportSave. It is recommended to gather SupportSave information as soon as you are able, since SupportSave captures all available debugging data; including SupportShow information.

For information about SupportShow CLI command, refer to the *Fabric OS Command Reference* manual.

Related CLI commands: **supportsave**, and **supportshow**. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

### 10.9.3 Port Error Statistics

To check port error statistics on a Director or Logical Switch, right-click the Director or Logical switch and select **Element Manager** → **Ports**, as shown in Figure 10-12.

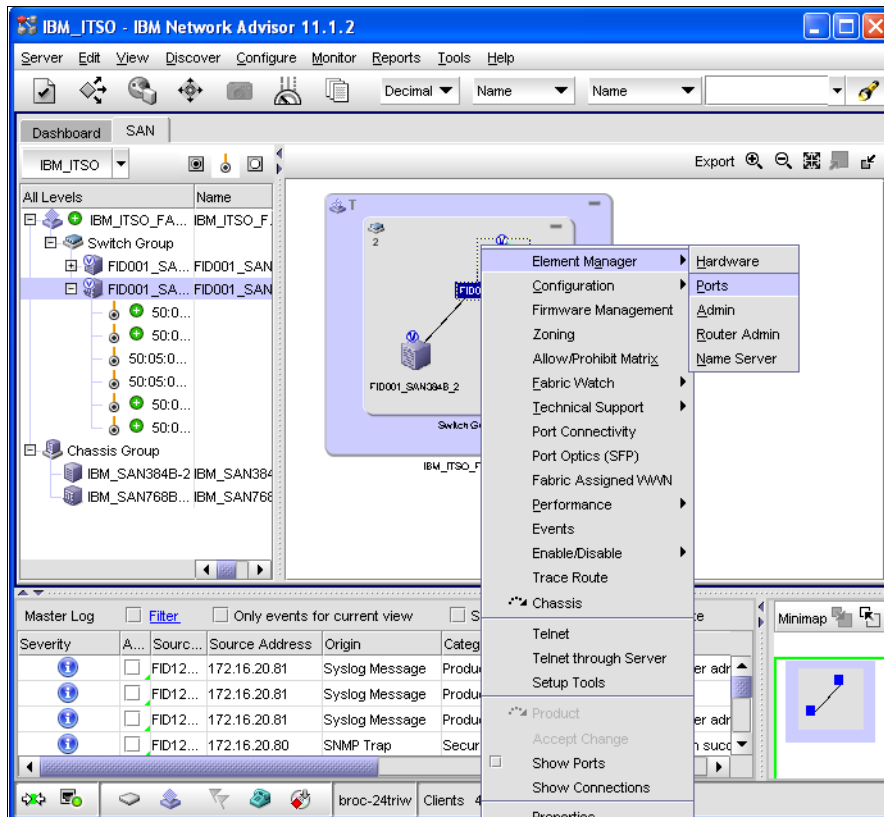


Figure 10-12 Selecting Port Statistics

After you are on the Port Administration panel, select the port you want to check and click **View Details**. Port Administration panel now shows you detailed port information. Keep the



Advanced view mode on the panel. To change the panel view mode, use the button at the top right of the panel, see the red box in Figure 10-13.




Figure 10-13 Port Administration - “Advanced mode” Detailed Port information

Select the **Port Statistics** tab. There are three main blocks inside this tab:

- ▶ Port Information: basic port data information.
- ▶ Port Statistics Configuration: this section allows us to configure the data statistics view, Absolute values (accumulated values) and Delta values (last 30 seconds, configured on Update Interval In Secs field). See blue box in Figure 10-13.
- ▶ Basic/Advanced/Error Details tabs: port statistics information. See the green box in Figure 10-13.

The Error Details tab gives us error statistics information to use during a Ficon problem determination, see Figure 10-14.




Figure 10-14 Error Details

**Note:** Use Delta view to check the last Update Interval seconds of port statistics. In Figure 10-13 on page 369 is the last 30 seconds. This option helps you with problem determination to be able to see the evolution of the errors.

Related CLI commands: **porterrshow**, **portstatshow**, and **portstatsclear**. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

## 10.9.4 Fibre Channel (FC) Troubleshooting

IBM Network Advisor has the following tools to help you analyze problems:

- ▶ FC Trace Route
- ▶ Device Connectivity
- ▶ Fabric Device Sharing
- ▶ Diagnostic Port Test

Select **Configure** → **FC Troubleshooting** to access these tools as shown in Figure 10-15.

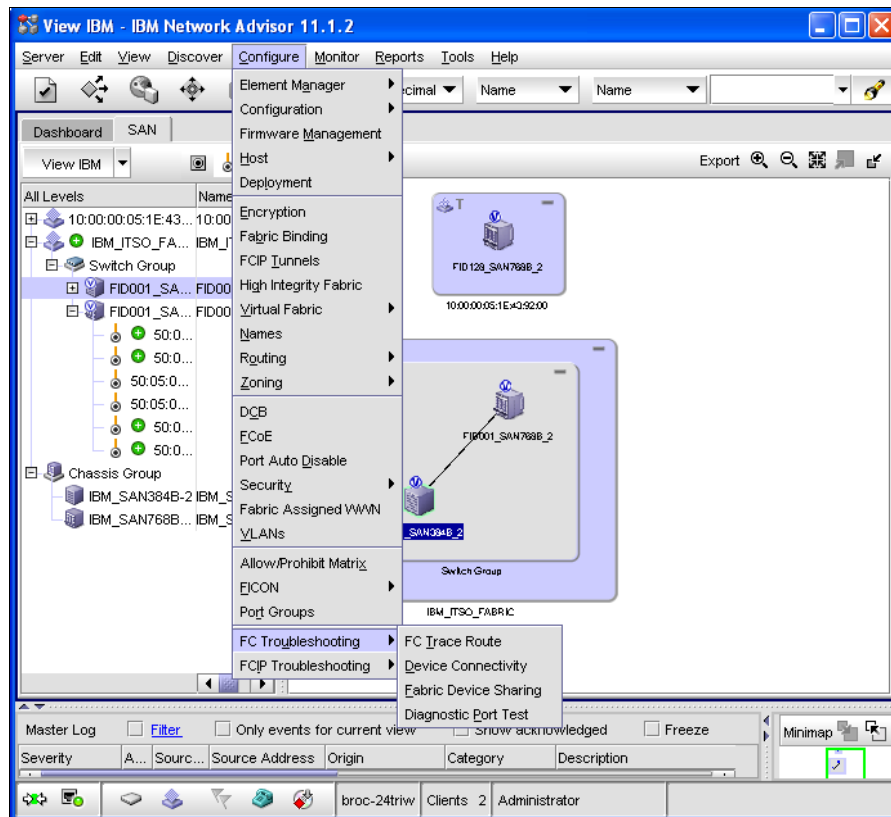


Figure 10-15 FC Troubleshooting menu

### FC Trace Route

FC Trace Route is used to trace the route used between two ports. This feature displays detailed routing information from a source port of the local switch to a destination port located on another switch.

**Tip:** That function is useful, for example, when you suspect you are having a problem on a E\_port connection. Trace Route helps you to locate the trunk used to communicate between the source and destination port. Having that information can significantly reduce the analysis range.

Select **Configure** → **FC Troubleshooting** → **FC Trace Route**, as shown in figure Figure 10-15.

Trace Route panel has the ability to specify the ports using the port FC address or by selecting the ports using the down side frames; see Figure 10-16.

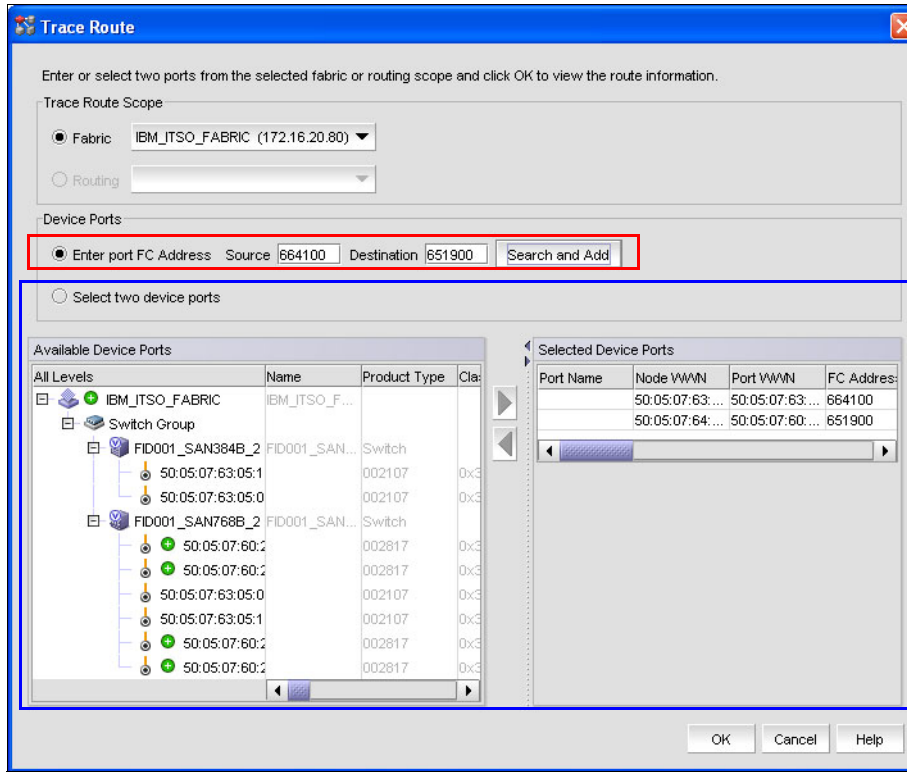


Figure 10-16 FC Trace Route panel

Routing information appears with the Trace Route Summary panel, see Figure 10-17 on page 372.

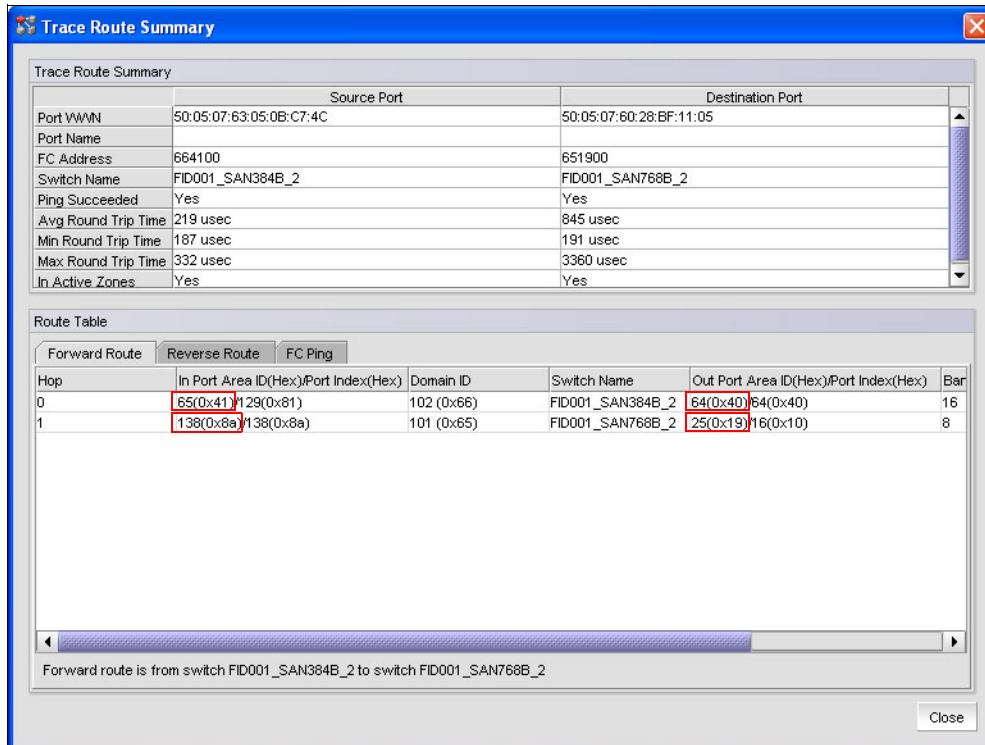


Figure 10-17 Trace Route Summary

At the bottom of Trace Route summary panel you have three tabs:

- ▶ Forward Route: Routing information to go from source port to destination port.
- ▶ Reverse Route: Routing information to go from destination port to source port.
- ▶ FC Ping: Ping test information to check ports communication health.

Related CLI commands: **urouteshow**, and **fcping**. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

## Device Connectivity

Device Connectivity performs diagnostics to a device connection.

The following diagnostic tests are performed:

- ▶ Device Status
- ▶ Switch port health status
- ▶ Zone configuration in the fabric

Select **Configure** → **FC Troubleshooting** → **Device Connectivity**, as shown in figure Figure 10-15 on page 370.

## Fabric Device Sharing

Fabric Device Sharing is a troubleshooting FC Routing that is not supported for FICON environment.

## Diagnostic Port Test

Diagnostic Port Test can be used for 16 Gbps capable E\_ports.

The following checks are performed:

- ▶ Electrical Loopback Test (E-Wrap)
- ▶ Optical Loopback Test (O-Wrap)
- ▶ Measure link distance
- ▶ Link Saturation Test

**Note:** Diagnostic Port Test is only available for E\_ports with 16 Gb capable (2499-816 and 2499-416). Diagnostic Port Test used on a 16 Gb port capable with 10 Gb optic, does not have full capability: E-Wrap test and O-Wrap test are not performed, and the cable distance test accuracy is within 50 meters (being 5 meters on 16 Gb optic port). Link saturation is performed for both, 16 Gb optics and 10 Gb optics.

Select **Configure** → **FC Troubleshooting** → **Diagnostic Port Test**, as shown in figure Figure 10-15 on page 370.

On the Diagnostic Port Test panel, select the E\_ports to which you want to perform the diagnostics by using the arrow buttons. See Figure 10-18.

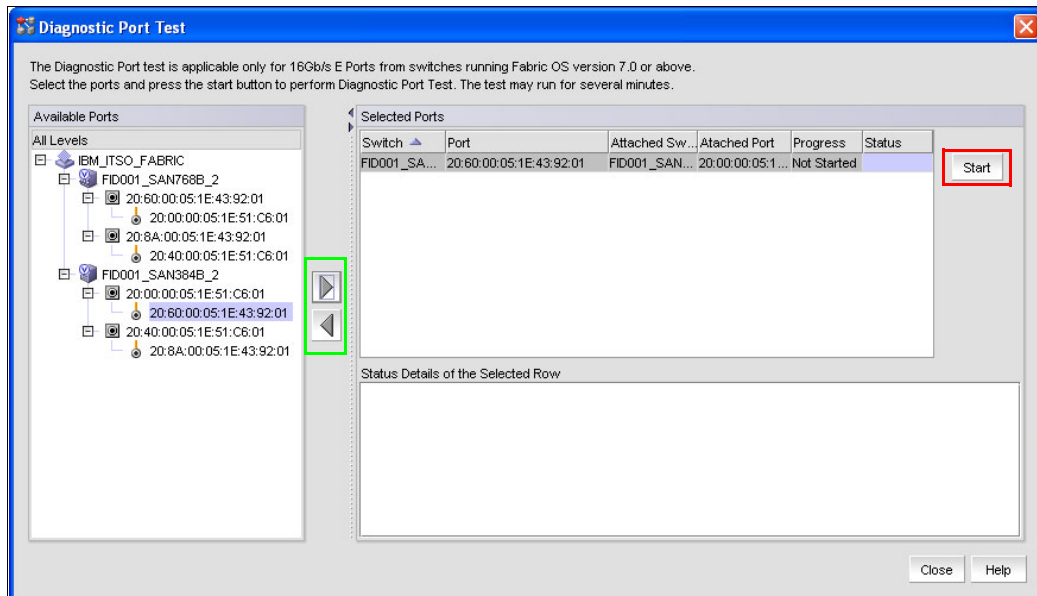


Figure 10-18 Diagnostic Port Test panel

After selecting the ports and pressing the Start button to start the diagnostics, a Warning panel appears to advise you about the consequences of performing a Diagnostic Port Test. You can click **OK** to continue or **Cancel** to abort the diagnostics. See Figure 10-19 on page 374.



Figure 10-19 Diagnostic Port Warning

**Tip:** Performing a Diagnostic Port Test affects the selected port and its attached ports. To avoid I/O interruptions it is highly recommended to isolate the affected ports.

One procedure can be:

- ▶ Perform a Port Decommission to the affected E\_ports  
The port you want to perform the diagnostics on and its attached E\_port on the remote switch. Port Decommission disables the port without generating an I/O interruption.
- ▶ Move the E\_port and its remote attached E\_port to another Logical Switch  
This is to ensure that after enabling the E\_port to perform the Diagnostic Port, the E\_port is not going to get used by I/O data traffic.
- ▶ Enable the E\_port and its remote attached E\_port  
After you have moved the E\_ports to another Logical Switch, you have to enable the port to be available to perform the Diagnostic Port. Otherwise, the port will not appear on the Diagnostic Port Panel.
- ▶ Perform Diagnostics Port

After you have done the Diagnostics, you can return the E\_port and its attached E\_port to their original Logical Switch.

The Progress column shows test status: Not started, in progress, or completed. The Status column shows us the overall diagnostics result: Success or Failed. See Figure 10-20 on page 375.

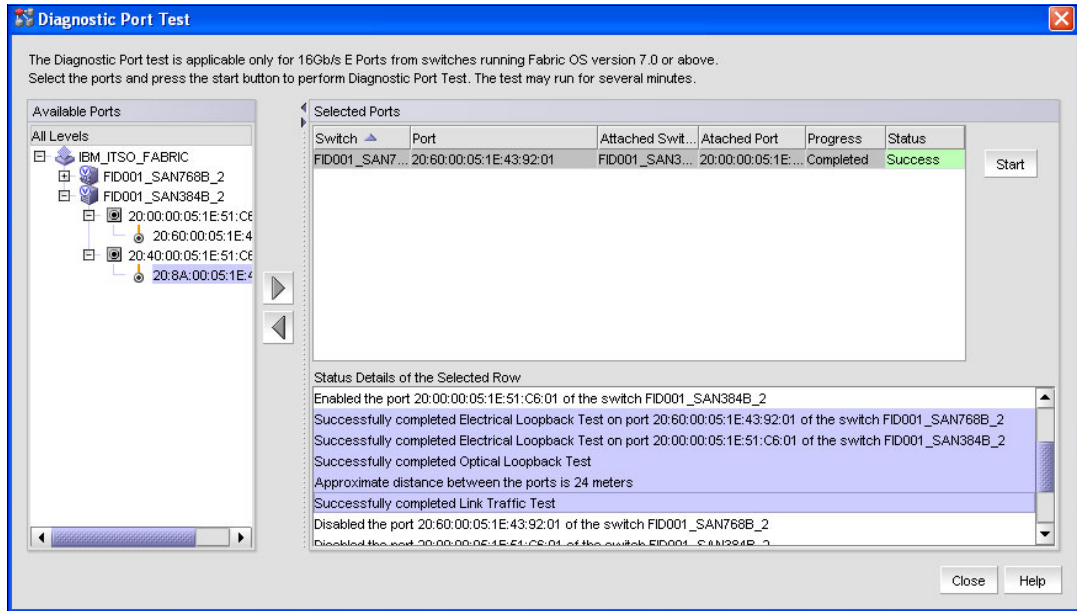


Figure 10-20 Diagnostic Port Test resulting information

**Tip:** The Diagnostic Port Test is especially useful to pre-test cables and optics before configuring online new ISL links: individual E\_ports or Trunk members. You can combine this function with `fcping` CLI command.

Related CLI commands: `portcfgdport`, `portdporttest`, and `portdecom`. See 10.9.5, “Helpful Command Line Interface commands” on page 375.

## 10.9.5 Helpful Command Line Interface commands

The following Command Line Interface (CLI) commands can be useful during Problem Determination. Although most of them can be performed on IBM Network Analysis application, it is helpful to have related CLI command.

**Note:** We do not have a direct CLI command for all IBM Network Advisor features, but, in most of the cases, you can find an alternative.

For detailed information about CLI commands, refer to the *Fabric OS Command Reference* manual, which can be found at:

<http://www.mybrocade.com>

The helpful CLI commands are:

- ▶ `auditcfg` modifies and displays the audit log filter configuration. This command allows you to set filters by configuring certain classes, to add or remove any of the classes in the filter list, to set severity levels for audit messages, and to enable or disable audit filters.
- ▶ `auditdump` displays or clears the audit log on the switch.
- ▶ `errdump` prints all error messages for the logical switch.

- ▶ **supportshow** displays switch information for debugging purposes. It is possible to perform the display by port, or range of ports, and you can also perform the display using defined groups for specific information to collect.
- ▶ **supportsave** collects debugging data. Supportsave gathers RASLOG, TRACE, supportshow, core file, FFDC data, and other support information.
- ▶ **porterrshow** displays an error summary of all ports with one line per port.
- ▶ **portstatsclear** lets you clear port hardware statistics.
- ▶ **portstatsshow** displays port hardware statistics counters.
- ▶ **urouteshow** displays routing information between active ports.
- ▶ **ficonshow** displays the contents of a FICON database. Using RNID operand, the command shows you the device node identification database.
- ▶ **portdecom** allows you to safely remove an E\_port. This automated mechanism gives frames in flight the possibility to complete.
- ▶ **portcfgdport** configures a diagnostic port (D\_port). A D\_port does not carry any interswitch traffic or data traffic. The purpose of this type of port is running link-level diagnostics between two switches.
- ▶ **portdporttest** manually starts or finishes a Diagnostic Port Test on a D\_port. The port must be configured as a D\_port and physically connected to a second D\_port on a remote switch.
- ▶ **fcping** is used to send a Fibre Channel ELS Echo request between two ports (source and destination), to a single device, or execute a SuperPing. This command has several useful configuration parameters, some of them are: number of frames, length size, and interval wait.

**Note:** The fcping command is a helpful PD command that allows you to perform specific link testing. It is not available as a IBM Network Advisor feature.

- ▶ **bottleneckmon** configures the parameters to monitor and report latency and congestion bottlenecks on F\_ports and E\_ports.

**Tip:** You can use *help* CLI command to display information related to CLI commands. It is possible to use the *grep* operand to improve your command searching.

For example: Displaying all commands with *fcp*

```
FID001_SAN768B_2:FID1:admin> help |grep fcp
fcping                Sends a FC ELS ECHO request to a port
fcplogclear           Clear the FCP log
fcplogdisable         Disable logging of FCP events
fcplogenable          Enable logging of FCP events
fcplogshow            Display FCP probing log information
fcpprobeshow          Display FCP probing information
fcprlsshow            Display FCP RLS (read link state) information
FID001_SAN768B_2:FID1:admin>
```

Now, you can use, for example, the **help fcping** command to display detailed contents of fcping CLI command.



## 10.10 HMC PD Tools

Refer to Appendix E, “Using HMC and SE for problem determination information” on page 425, for detailed explanation on the use of HMC and SE panels to obtain information when performing problem determination tasks.

## 10.11 Common z/OS FICON error message: Interpretation

We describe the most common z/OS FICON messages (IOS) and their respective meanings, explanations, and required actions when applicable. You might receive these messages when encountering these types of problems:

- ▶ Failure to complete initialization or establish a logical path
- ▶ Failure detected in communications of command and data by either the channel or the control unit

**Note:** The messages covered here might change, depending on the z/OS release and version. APARs can also implement architectural changes regarding how messages are generated and displayed. We recommend that readers refer to the current z/OS Messages library for reviewing the output whenever a discrepancy is noted.

Refer to the *z/OS MVS System Messages*, SA22-7637, for specific message details and actions.

You can also find information about messages using the IBM online facility *LookAt*. You can find z/OS messages explanations in a easy and fast way.

<http://www-03.ibm.com/systems/z/os/zos/bkserv/lookat/index.html>

### **IOS001E**

devn, INOPERATIVE PATH chp[,chp,chp...]

Message indicates that one or more paths, that was previously operational (Path Operational Mask (POM) was ON), went inoperative. This is indicated by the Path Not Operational Mask (PNOM) Unit Control Word (UCW) field. For more information about PNOM and POM fields, refer to “Displaying the Analyze Subchannel Information PD panel” on page 442.

The message text shows:

**devn** Device in use by an I/O operation when the failure occurred.

**chp** Path in use by the device that become inoperative.

**Tip:** Check for a disabled Control Unit interface, an altered FICON director configuration, or a disabled channel path. Refer to the explanations of messages “IOS2001I” on page 386 or “IOS2002I” on page 388 in this section for possible reasons why you might receive message IOS001E.

### **IOS002A**

devn, NO PATHS AVAILABLE

STATUS FOR PATH(S) chp,chp,...

Indicates that the device did not respond on any online path.

Message text shows:

**devn** Device number selected for operation.

**chp** Channel Paths available to device access.

Possible reasons that can cause IOS002A message:

- ▶ There are no paths available to the device.
- ▶ During device initialization, the system encountered a permanent I/O error.

The I/O requester will be kept waiting until an I/O interrupt is received. You can do one of the following actions to simulate an I/O interrupt and cause queued requests to be retried:

- ▶ Enter a VARY DEVICE ONLINE command for device: **VARY devn,ONLINE**
- ▶ Enter a VARY PATH ONLINE command for a path to the device: VARY PATH(devn,chp),ONLINE
- ▶ If a serial fiber is not connected, connect the serial fiber to device devn

**Note:** On all paths that present Deferred Condition Code 3 (DCC=3), Dynamic Pathing Switch Validation returns RC=12, to indicate to the caller that the “no paths” bit is to be turned on in the device UCB and wait for the CU to become operational, generating a State-Change Ready interrupt.

Power drop of the CU or FICON Director can be the cause of a DCC=3.

## ***IOS050I***

**CHANNEL DETECTED ERROR ON devn,chp,cmd,stat [,PCHID=pppp]  
[,EXTSS=ss]**

Indicates that a Channel Error was detected: Interface Control Check (IFCC), Channel Control Check (CCC), Channel Data Check (CDC),...

The message text shows:

**devn** Device in use by an I/O operation when the failure occurred.

**chp** Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks.

**cmd** Failing command being executed by the device, if known. Otherwise, this field will be set to asterisks.

**stat** Device and subchannel status, if known. Otherwise, this field will be set to asterisks.

**PCHID=pppp** System z server Physical Channel Identifier.

**EXTSS=ss** Subchannel extended status.

**Tip:** Run EREP or examine the CHPID work area on SE to look for the IFCC Threshold Exceeded CHPID status. Alternatively, use the “IFCC and Other Errors” service task on the SE to further investigate possible reasons why this message was issued.

Refer to E.4, “Displaying FICON IFCC and Other errors” on page 453.

**Note:** Introduced in zOS 1.13, IOS050I and IOS051I have a new message extension that informs you about the component that has detected the error. Refer to “IOS054I” on page 380 for more information.

### ***IOS051I***

**INTERFACE TIMEOUT DETECTED ON devn,chp,cmd,stat [,PCHID=pppp]  
[,EXTSS=ss]**

Channel subsystem detected an interface timeout IFCC.

The message text shows:

|                   |   |
|-------------------|---|
| <b>devn</b>       | Device in use by an I/O operation when the failure occurred.  |
| <b>chp</b>        | Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks.                  |
| <b>cmd</b>        | Failing command being executed by the device, if known. Otherwise, this field will be set to asterisks (*). |
| <b>stat</b>       | Device and subchannel status, if known. Otherwise, this field will be set to asterisks.                     |
| <b>PCHID=pppp</b> | System z server Physical Channel Identifier.  |
| <b>EXTSS=ss</b>   | Subchannel extended status.   |

**Tip:** IOS component checks the state of the timeout bit before generating the IOS051I message for FICON channels. The timeout indication is reported by the channel in the Extended Status Word Timeout (ESWT) bit. If timeout bit is not on, then IOS050I will be issued instead.

**Note:** Introduced in zOS 1.13, IOS050I and IOS051I have a new message extension that informs you about the component that has detected the error. Refer to “IOS054I” on page 380 for more information.

### ***IOS052I***

**CHANNEL EVENT REPORTED FOR devn,chp,cmd,stat [,PCHID=pppp]  
[,EXTSS=ss]**

Indicates that the Channel detected an error (e.g., cross check errors), while the system was operating a device, and the error has automatically been recovered by the channel.

The message text shows:

|             |   |
|-------------|---|
| <b>devn</b> | Device in use by an I/O operation when the failure occurred.  |
| <b>chp</b>  | Channel path identifier (CHPID), if known. Otherwise, this field will be set to asterisks.              |
| <b>cmd</b>  | Failing command being executed by the device, if known. Otherwise, this field will be set to asterisks. |
| <b>stat</b> | Device and subchannel status, if known. Otherwise, this field will be set to asterisks.                 |

**PCHID=pppp** System z server Physical Channel Identifier.

**EXTSS=ss** Subchannel extended status.

**Note:** As the error has automatically been recovered by the channel, the message is issued to the SYSLOG only to provide an audit trail for serviceability. No operator console message appears.

### ***IOS054I***

sdddd,chnp ERRORS DETECTED BY component[, component]...

Introduced with z/OS 1.13 as a message extension of IOS050I and IOS051I. The message provides the hardware component(s) that detected the error.

Message text shows:

**sdddd** Subchannel set identifier and the device in use by an I/O operation when the failure occurred.

**chnp** Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks.

**component** Detecting component and can be one or more of the following:

- ▶ CHANNEL
- ▶ CHAN SWITCH PORT
- ▶ CU SWITCH PORT
- ▶ CONTROL UNIT

**Tip:** Run EREP to look for more information:

- ▶ SLH EDIT report - Extended Subchannel Logout Data: information about errors detected (Link Failure, Loss of Synch, Loss of Signal,...). See 10.6, "FICON Purge Path Extended" on page 357.
- ▶ Link Maintenance Information Detail Edit (A2) report: possible link errors detected. This record is associated with SYSLOG messages IOS580E and IOS581E.

**Note:** This message will always be preceded by messages IOS050I or IOS051I identifying the hardware component that detected the error. The message appears if a channel error has occurred. You can use the detecting component information to isolate the faulty hardware component.

### ***IOS071I***

devn,chnp,jobname,text [STATUS: statustext]

System finds the condition appearing in the message *text* and one of the following occurred:

- ▶ The condition existed for a longer time than the acceptable time defined by the installation.
- ▶ A Missing Interrupt Handler (MIH) exit routine indicates that the system will not reset the device with a Clear Subchannel (CSCH) instruction.

Message text shows:

**devn** Device in use by an I/O operation when the failure occurred.

**chnp** Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks.

**jobname** Name of the job associated with the device.  
**text** Description of the missing interruption condition.  
**statustext** State of the I/O operation prior to detecting the missing interruption condition.

Condition appearing in the message text can be one of the following issues:

- ▶ **MISSING CHANNEL AND DEVICE END** or **MISSING DEVICE END**  
A hardware error occurred.
- ▶ **START PENDING**  
Hardware error occurred and the channel has an I/O request to process. The device is not active or it is busy on another channel path interface. In case the device is reserved by another system, message IOS431I might follow identifying the system holding the reserve.
- ▶ **IDLE WITH WORK QUEUED**  
A hardware or software error occurred. The channel has no active I/O requests.
- ▶ **HALT SUBCHANNEL INTERRUPT MISSING** or **CLEAR SUBCHANNEL INTERRUPT MISSING**  
Hardware error occurred.
- ▶ **HYPERSWAP INITIATED TIMEOUT**  
A Hyperswap was in progress. MIH detects this condition and immediately times out any active I/O for any device being swapped. The I/O request is cancelled and re queued for later execution after the swap is complete.

**Tips:**

- ▶ The jobname can be \*MASTER\*. In this case, IOS most likely is recovering I/O such as Unconditional Reserve (U/R) or Dynamic Path Selection (DPS) and the operation is generating the MIH condition.
- ▶ If Missing CE/DE for CTC devices, probably the system on the other end was not properly running (in a loop or wait condition, busy with I/O interrupts being disabled, or in a thrashing state for paging I/Os). Look at the RMF report for the paging I/O rate.
- ▶ An IDLE WITH WORK QUEUED message might occur if an Immediate Execution (IMEX) Request is queued while a sense operation is pending for a device. This occurs because the IMEX request cancels the pending sense operation so it can execute first. However, IOS will not be able to start the IMEX request because the sense I/O is pending. Thus, MIH will find this condition and declare it as an IDLE WITH WORK QUEUED condition.
- ▶ For Start-Pending, Channel Path chp might not be correct. It is just the previous successfully selected path. To identify the correct path, refer to the EREP MIH record.

**IOS075E**

devn, RECURRING MIH CONDITION FOR THIS DEVICE

System found a recurring missing interruption for a device. Recursive MIH condition was not cleared by the previous MIH Clear Subchannel (CSCH) recovery sequence.

Message text shows:

**devn** Device in use by an I/O operation when the failure occurred.

Look for more information about previous messages IOS071I, IOS077E, or IOS076E. The system also writes a LOGREC data set error record.

**Tips:**

- ▶ If you specify an MIH HALT interval that is too short (such as 1 second), change it to at least 2 seconds because the interval of the FICON Interface-Timeout IFCC (No Response from CU) is 2 seconds.
- ▶ CSCH was rescheduled due to the recursive MIH. This means that the recursive MIH condition was detected before the original MIH recovery completed. As a result, the original MIH message might not have been issued and this message was issued instead.
- ▶ The difference from message IOS077E is that this message is issued (instead of IOS077E) when the MIH message type is unknown (whether IOTIMING, Primary Status Missing, Secondary Status Missing, Start Pending, Halt Interrupt Missing, or Clear Interrupt Missing), as indicated by the CONDITON flag being OFF in the MIH message module.
- ▶ CSCH might not complete quickly if the CSS-selected CHPID is in a permanent busy condition based on the Last Path Used Mask (LPUM) of the subchannel.

**IOS076E**

devn,chnp,jobname,text [,STATUS:statustext]

System finds the condition appearing in the message *text* and one of the following occurred:

- ▶ The condition existed for a longer time than the acceptable time defined by the installation.
- ▶ A Clear Subchannel (CSCH) interruption is missing.

Message text shows:

|                   |  |
|-------------------|--|
| <b>devn</b>       | Device in use by an I/O operation when the failure occurred.                               |
| <b>chnp</b>       | Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks. |
| <b>jobname</b>    | Name of the job associated with the device.  |
| <b>text</b>       | Description of the missing interruption condition.   |
| <b>statustext</b> | State of the I/O operation prior to detecting the missing interruption condition.          |

The message indicates missing interrupt conditions where the MIH recovery did not reset the device (that is, it did not issue the CSCH); for example:

- ▶ The MIH was for CLEAR Subchannel interrupt missing.
- ▶ Or the TAPE MIH exit requested no CSCH to be issued when the current active I/O contains a long-running CCW, such as Forward-Space File and so on.

For a DASD device, there is no case where this occurs.

**Tip:** This message is issued when there is an active I/O operation that was not cleared by the CSCH recovery action sequence that was performed.

## ***IOS077E***

devn, chp, jobname, text

System found a condition that existed for longer time than the acceptable time defined by the installation.

Message text shows:

**devn** Device in use by an I/O operation when the failure occurred.  
**chp** Channel Path Identifier (CHPID), if known. Otherwise, this field will be set to asterisks.  
**jobname** Name of the job associated with the device.  
**text** Description of the missing interruption condition.

**Tip:** This message means that another missing interrupt condition (MIH) was detected before the previous MIH recovery completed. Look for a previous IOS071I message.

## ***IOS100I***

DEVICE devn BOXED, [LAST PATH yy LOST,] [text]

An error occurred on the last channel path to a device.

Message text shows:

**devn** Device in use.  
**yy** Last path available to *devn* that has been lost.  
**text** Condition detected.

This message indicates that one of these conditions was detected, where text is one of the following:

- ▶ RESERVE LOST
- ▶ CANNOT RE-RESERVE
- ▶ STILL RESERVED
- ▶ RESERVE MAY BE LOST
- ▶ MAY BE RESERVED
- ▶ ASSIGN LOST
- ▶ CANNOT RE-ASSIGN
- ▶ STILL ASSIGNED
- ▶ SELF-DESCRIPTION RECOVERY

**Tip:** In the Channel Path recovery, the last path (CHPID) to the device has been varied offline. The device has been Boxed. Any new I/O request to this device will result in Permanent I/O Errors.

If other channel paths are available to the device, vary them online. This condition is generally associated to a hardware failure in the CU or device.

## ***IOS101I***

DEVICE devn {FORCED OFFLINE | BOXED,} {text}

An error occurred on the last channel path to a device.

Message text shows:

**devn**            Device in use.  
**text**            Condition detected.

This message indicates that one of these conditions was detected, where *text* is one of the following:

- ▶ RESERVE LOST
- ▶ CANNOT RE-RESERVE
- ▶ RESERVED ON FAILING PATH *pp*
- ▶ RESERVE MAY BE LOST
- ▶ MAY BE RESERVED ON PATH *pp*
- ▶ ASSIGN LOST
- ▶ CANNOT RE-ASSIGN
- ▶ WAS RESERVED
- ▶ WAS ASSIGNED
- ▶ STILL ASSIGNED
- ▶ STILL RESERVED
- ▶ ASSIGNED IN FAILING PATH *pp*

**Tip:** The Device Recovery routine invoked in the Channel Path Recovery detected one of these conditions as it tried to Re-Reserve/Assign the device after RCHP - Reset Channel Path (System Reset). Any new I/O request to the device will result in a Permanent I/O error.

This condition is generally associated to a hardware failure in the CU or device.

This condition can result from parameter BOX\_LP specified for HOTIO or TERMINAL statements in the IECIOSxx PARMLIB member. Note that BOX\_LP=(ALL) is the default value.

## ***IOS102I***

DEVICE devn {BOXED | FORCED OFFLINE},{text}

System boxed a device or forced the device offline.

Message text shows:

**devn**            Device number.  
**text**            Condition detected.

This message indicates that one of these conditions was detected, where *text* can be one of the following:

- ▶ OPERATOR REQUEST  
Operator can be issued: VARY devn,OFFLINE,FORCE command.
- ▶ PERMANENT ERROR
- ▶ C.U.I.R. REQUEST  
Control Unit Initiated Recovery (C.U.I.R.) has received a request to fence a device or the last path to the device.



- ▶ WAS RESERVED
- ▶ STILL RESERVED
- ▶ MAY BE RESERVED
- ▶ WAS ASSIGNED
- ▶ STILL ASSIGNED
- ▶ ASSIGN LOST
- ▶ NO PATHS

Online paths to a device did not respond when the operator tried to validate the status of a path.

- ▶ RESERVE LOST
- ▶ MOVE ALIAS FAILURE

WLM has requested that a Parallel Access Volume (PAV) Alias device be moved from one base device to another. Alias movement has failed repeatedly. IOS boxes the alias to prevent WLM from continuously selecting this same Alias to be moved.

- ▶ DEFERRED BOX PROCESSING COMPLETE

The device had its BOX processing deferred and this message is issued to indicate that deferred BOX processing is now complete and the device is now in the boxed state. See “IOS071I” on page 380 for more information.

This message itself might not show the reason of the device BOX. It is usually accompanied by another IOS message (for example, IOS451I or IOS109E and so on) that might contain the reason for the device being boxed. Also, be aware that if these messages are issued within the same time window (milliseconds), the SYSLOG message order might sometimes be out of sequence (reversed).

**Tip:** A permanent error on all Paths, or Reserve/Assign Lost in DPS Array, was detected during DPS Validation with message IOS451I.

## ***IOS107I***

### **DEVICE devn BOX PROCESSING DEFFERRED**

Indicates that the system deferred the device devn box processing to allow recovery processing to run.

Message text shows:

**devn**            Device number.

**Tip:** Certain recovery processing, such as IBM HyperSwap® operations, might be able to provide continuous availability to the device by swapping the faulty device (that is, the primary device) with a fully operational “backup” device (that is, the secondary device) prior to permanently boxing the faulty device.

At the end of this process, an IOS102I message will indicate that the boxing process has completed.

## ***IOS162A***

### **CHPID chp ALERT, UNSOLICITED MALFUNCTION INTERRUPT**

While trying to service a request from a device, the channel subsystem found an error before the system determined the device number.

Message text shows:

**chp** Channel Path Identifier (CHPID).

### ***IOS163A***

CHPID chp ALERT, NO ASSOCIATED SUBCHANNEL FOR DEVICE

Hardware error occurred while the system was trying to service a request from a device.

Message text shows:

**chp** Channel Path Identifier (CHPID).

### ***IOS167I***

DEVICE devn RESTORED BY SUBCHANNEL RECOVERY

Device subchannel that has undergone recovery and has been restored to an operational state.

Message text shows:

**devn** Device number.

**Note:** This message is introduced with z/OS 1.13.

### ***IOS2001I***

devn,INOPERATIVE PATHS chp,chp,...

STATUS FOR PATH(S) chp,chp,...

errortype

[errormsg]

Device was unable to operate on the specified channel paths. This message appears only in the SYSLOG.

Message text shows:

**devn** Device number.

**chp** Inoperative Channel Path Identifier (CHPID).

**errortype** Reason for the inoperative path.

**errormsg** Reason for the inoperative path.

When such a condition is present, the result of the display matrix device command, D M=DEV(devn,(chp)), will report it.

When message MSG IOS001E, MSG IOS002A, or MSG IOS450E is issued, IOS will retrieve the not-operational reason from the Channel Subsystem and issue message IOS2001I (for

MSG IOS001E or MSG IOS450E) or MSG IOS2002I (for MSG IOS002A), to SYSLOG only, with text describing the reason.

**Tip:** The message displays the STATUS FOR PATH(S) chp, chp, . . . not-operational reason text. The second line of text, if it exists, contains the specific reason for the not-operative path.

The paths specified were found to be not operational for the specified device. Some FICON-relevant Status Code message texts are listed here:

- ▶ 50: CHANNEL SUBSYSTEM DETECTED A LINK FAILURE CONDITION):
  - 10: LOSS OF SIGNAL OR SYNCHRONIZATION CONDITION RECOGNIZED
  - 20: NOT OPERATIONAL SEQUENCE RECOGNIZED
  - 30: SEQUENCE TIMEOUT RECOGNIZED
  - 40: ILLEGAL SEQUENCE RECOGNIZED
- ▶ 60: CHANNEL LINK LEVEL FACILITY IN OFFLINE RECEPTION  
This is the case when the Channel-to-CU logical path (H/W link-level) has been de-established (broken).
- ▶ 70: PORT REJECT WAS ENCOUNTERED:
  - 10: ADDRESS INVALID ERROR
  - 11: UNDEFINED DESTINATION ADDRESS ERROR
  - 12: DESTINATION PORT MALFUNCTION
  - 13: DYNAMIC SWITCH PORT INTERVENTION REQUIRED
- ▶ 80: LINK LEVEL REJECT WAS ENCOUNTERED:
  - 01: TRANSMISSION ERROR
  - 05: DESTINATION ADDRESS INVALID ERROR
  - 07: RESERVED FIELD ERROR
  - 08: UNRECOGNIZED LINK CONTROL FUNCTION
  - 09: PROTOCOL ERROR
  - 0A: ACQUIRE LINK ADDRESS ERROR
  - 0B: UNRECOGNIZED DEVICE LEVEL
- ▶ 90: CHANNEL LINK LEVEL FACILITY DETECTED A LINK ERROR:
  - 10: CONNECTION ERROR
  - 20: TRANSMISSION ERROR
  - 30: PROTOCOL ERROR
  - 40: DESTINATION ADDRESS INVALID ERROR
  - 50: DEVICE LEVEL ERROR (50)
  - 60: LINK RCYVY THRESHOLD EXCEEDED FOR ALL LOGICAL PATHS OF DEST LINKS.

**Note:** For System z server FICON channels, this is the case where the CU logical path connection has been fenced because the up/down (flapping) error count has been exceeded. (Channel Path or Link level threshold exceeded). You must enter the V PATH (dddd,pp),ONLINE command to issue a Reset Link Recovery Threshold CHSC command, or Config CHPID OFF then ON (using a CF z/OS command or via the HMC).

- ▶ A0: LOGICAL PATH IS REMOVED OR NOT ESTABLISHED:
  - 01: PACING PARAMETERS ERROR
  - 02: NO RESOURCES AVAILABLE
  - 04: DESIGNATED CONTROL UNIT IMAGE DOES NOT EXIST
  - 05: LOGICAL PATH PRECLUDED BY CONFIGURATION AT CONTROL UNIT IMAGE
  - 06: LINK RECOVERY THRESHOLD EXCEEDED FOR LOGICAL PATH

**Note:** For System z servers FICON channels, this is the case where CU logical path connection has been fenced due to up/down (flapping) error count exceeded. (Channel Path or Link level threshold exceeded). You must enter the V PATH (dddd,pp),ONLINE command to issue a Reset Link Recovery Threshold CHSC command, or Config CHPID OFF then ON (using a CF z/OS command or through the HMC).

- ▶ B0: IN PROCESS OF INITIALIZING PATH:
  - 10: CONTROL UNIT DEVICE LEVEL INITIALIZATION IS NOT COMPLETE
  - 20: LINK BUSY CONDITION LAST ENCOUNTERED
  - 30: PORT BUSY CONDITION LAST ENCOUNTERED
  - 30: PORT BUSY CONDITION LAST ENCOUNTERED
- ▶ C0: SUBCHANNEL PATH AVAILABLE, BUT DEVICE NOT OPERATIONAL:
  - 30: PORT BUSY CONDITION LAST ENCOUNTERED
  - FF: NO FURTHER INFORMATION AVAILABLE OR UNKNOWN CONDITION

### ***IOS2002I***

devn,INOPERATIVE PATHS chp,chp,...

STATUS FOR PATH(S) chp,chp,...

errortype

[errormsg]

System selected a device for operation. The device did not respond to any online path. The system requires an online interrupt to retry I/O requests to the device. Message is written to SYSLOG only.

Message text shows:

|                  |  |
|------------------|--|
| <b>devn</b>      | Device number.                               |
| <b>chp</b>       | Inoperative Channel Path Identifier (CHPID). |
| <b>errortype</b> | Reason for the inoperative path.             |
| <b>errormsg</b>  | Reason for the inoperative path.             |

Refer to “IOS2001I” on page 386 for further information about errortype and errormsg.

**Tips:** All paths to the specified device were found not operational. Path status code is created by Path State Info routine issuing Store Subchannel Path Information (SSPI).

Note the following examples:

- ▶ FICON cables were disconnected at the CU side through FICON SW: STATUS FOR PATH(S) pp IN PROCESS OF INITIALIZING PATH (B0) NO FURTHER INFORMATION AVAILABLE OR UNKNOWN CONDITION (FF)
- ▶ FICON Director Power-Down (by power button): CHANNEL LINK LEVEL FACILITY IN OFFLINE RECEPTION STATE (60)

### ***IOS202E***

CHANNEL PATH chp FORCED OFFLINE - DEVICE IS: {devn | UNKNOWN}

System forces the channel path offline due to a permanent hardware error. External intervention is required to recover the channel path.

This message follows message IOS109E when CHP,F is the automatic recovery action for the HOTIO condition, or when CHP,K initiated unsuccessful channel path recovery.

Message text shows:

**chp** Channel Path Identifier (CHPID).  
**devn** Device number.  
**UNKNOWN** Appears if the device is unknown or no device is involved.

### ***IOS203I***

CHANNEL PATH chp SUCCESSFULLY RECOVERED - DEVICE IS: {devn | UNKNOWN}

An error occurred on a channel path to the specified device. One of the following occurred:

- ▶ Interface was either hung or broken.
- ▶ Device indicated that a reset event occurred.

This message follows message IOS109E when CHP,K is the automatic recovery action for the HOTIO condition, and channel path recovery is successful. This message follows message IOS207I when the system invokes control unit recovery.

System continues processing.

Message text shows:

**chp** Channel Path Identifier (CHPID).  
**devn** Device number.  
**UNKNOWN** Appears if the device is unknown or no device is involved.

### ***IOS210I***

PATH RECOVERY INITIATED FOR PATH chp ON CU cccc, REASON=rsntext

System displays this message when a path recovery-related error occurs and the installation has configured PATH\_SCOPE=CU option in the IECIOSxx PARMLIB member or using the

SETIOS command. The system tries to vary the channel path offline for all devices on the control unit.

Message text shows:

|                |   |
|----------------|---|
| <b>chp</b>     | Channel Path Identifier (CHPID).                                      |
| <b>cccc</b>    | Control unit for the device that encountered the path recovery error. |
| <b>rsntext</b> | Reason why the channel path is being varied offline.                  |

rsntext gives us the reason why the channel path is being varied offline:

► LINK RECOVERY THRESHOLD REACHED

Hardware FICON link recovery threshold has been reached and the path is no longer available to all devices on the control unit.

**Note:** This messages might be preceded by messages IOS2001I or IOS2002I.

► PATH ERROR THRESHOLD REACHED

The number of errors on a path over a period of time has reached an installation-specified threshold and the path needs to be taken offline to all devices on the control unit.

**Note:** This messages might be preceded by messages IOS050I or IOS051I.

► DYNAMIC PATHING ERROR

The system encountered an error on a path while validating a dynamic pathing device. As a result of the error, the path was taken offline from the device.

**Note:** This messages might be preceded by message IOS450E.

► REQUESTED BY DEVICE ERP ROUTINE

Device support error recovery procedure routine requested that control unit path error recovery be performed.

After the problem is corrected, you must bring the path online manually. The path is not automatically varied back online. You can use the following commands: VARY CU, VARY PATH, VARY device, or CONFIG CHP.

It is recommended that you first vary the path online to a single device to ensure that the problem has been corrected before varying the path online to the remaining devices.

**Note:** This message is introduced on z/OS 1.13 related to the new IECIOSxx configuration statement: PATH RECOVERY. You can find related information about *z/OS MVS Initialization and Tuning Reference* manual, SA22-7592, and *z/OS MVS System commands* manual, SA22-7627.

## ***IOS288A***

SYSTEM-INITIATED {ONLINE |OFFLINE} RECONFIGURATION IS IN PROGRESS FOR THE FOLLOWING CHPIDS:chp,chp-chp,chp,chp-chp... chp [WARNING: MESSAGE IS INCOMPLETE DUE TO SYSTEM ERROR.]

System is processing a reconfiguration request issued by the systems operator through the Hardware Management Console (HMC). The request is either an ONLINE or OFFLINE reconfiguration of the channel-paths specified by CHPID chp.

Message text shows:

**chp** Channel Path Identifier (CHPID).

*WARNING: MESSAGE IS INCOMPLETE DUE TO SYSTEM ERROR.*

This warning indicates that a system error occurred while the message was being constructed, resulting in an incomplete message.

**Note:** The system will try to re-configure the channel-paths specified in the message ONLINE or OFFLINE, as specified on the hardware management console. The system will issue message IOS289I after the reconfiguration is complete.

## ***IOS289I***

SYSTEM-INITIATED {ONLINE | OFFLINE} RECONFIGURATION HAS COMPLETED. [RECONFIGURATION WAS SUCCESSFUL FOR THE FOLLOWING CHPIDS: chp,chp-chp,chp,chp-chp...] [RECONFIGURATION FAILED FOR THE FOLLOWING CHPIDS: chp,chp-chp,chp,chp-chp...] [WARNING: MESSAGE IS INCOMPLETE DUE TO SYSTEM ERROR.]

This message reports the results of reconfiguration actions requested and reported by IOS288A, see "IOS288A" on page 391. System has completed processing a reconfiguration request issued by a system operator through the Hardware Management Console. The request is either an ONLINE or OFFLINE reconfiguration of the channel-paths specified by CHPID chp.

Message text shows:

**chp** Channel Path Identifier (CHPID):

- ▶ RECONFIGURATION WAS SUCCESSFUL FOR THE FOLLOWING CHPIDS:  
chp,chp-chp,chp,chp-chp... chp

Channel-paths specified in the message text were successfully reconfigured.

- ▶ RECONFIGURATION FAILED FOR THE FOLLOWING CHPIDS:  
chp,chp-chp,chp,chp-chp... chp

Reconfiguration failed for the channel-paths specified by the message text.

- ▶ WARNING: MESSAGE IS INCOMPLETE DUE TO SYSTEM ERROR.

Warning indicates that a system error occurred while the message was being constructed, resulting in an incomplete message.

### ***IOS580E***

LINK DEGRADED REPORTING CHPID=[chp | \*\*] {INCIDENT UNIT unit-ident  
IC=yy INCIDENT UNIT logical interface ident| INCIDENT UNIT DATA COULD  
NOT BE OBTAINED} [ATTACHED UNIT unit-ident] [ATTACHED UNIT logical  
interface ident]

Hardware error has occurred. A unit at one end of a optical interface link has detected a problem in the link or in one of the units attached to the link. As a result, data transfer across the link is degraded.

Message text shows:

- ▶ chp: Reporting Channel Path Identifier (CHPID). \*\* means the reporting source was not a channel path. The reporting CHPID is not always associated with the actual link incident and might just represent the CHPID that is used to report the link incident.
- ▶ INCIDENT UNIT unit-ident: Unit that detected the problem.
- ▶ yy: Incident code that identifies the type of incident that occurred.
- ▶ INCIDENT UNIT logical interface ident: Unit that detected the problem.
- ▶ INCIDENT UNIT DATA COULD NOT BE OBTAINED: Either the incident unit cannot be identified or it can be identified but the information might not be current and cannot be verified due to some problem, such as, loss of light.
- ▶ ATTACHED UNIT unit-ident: Unit connected to the other end of the link.
- ▶ ATTACHED UNIT logical interface ident: Unit connected to the other end of the link.

**Note:** Print LOGREC data set error record. See 10.5, “FICON link incident reporting” on page 357.

### ***IOS581E***

LINK DEGRADED REPORTING CHPID=[chp | \*\*] {INCIDENT UNIT unit-ident  
IC=yy INCIDENT UNIT logical interface ident| INCIDENT UNIT DATA COULD  
NOT BE OBTAINED} [ATTACHED UNIT unit-ident] [ATTACHED UNIT logical  
interface ident]

Hardware error has occurred. A unit at one end of a optical interface link has detected a problem in the link or in one of the units attached to the link. As a result, data transfer across the link has stopped.

Message text shows:

- ▶ chp: Reporting Channel Path Identifier (CHPID). \*\* means the reporting source was not a channel path. The reporting CHPID is not always associated with the actual link incident and might just represent the CHPID that is used to report the link incident.
- ▶ INCIDENT UNIT unit-ident: Unit that detected the problem.
- ▶ yy: Incident code that identifies the type of incident that occurred.
- ▶ INCIDENT UNIT logical interface ident: Unit that detected the problem.
- ▶ INCIDENT UNIT DATA COULD NOT BE OBTAINED: Either the incident unit cannot be identified or it might be identified but the information may not be current and cannot be verified due to some problem, such as, loss of light.



- ▶ ATTACHED UNIT unit-ident: Unit connected to the other end of the link.
- ▶ ATTACHED UNIT logical interface ident: Unit connected to the other end of the link.

**Note:** Print LOGREC data set error record. See 10.5, “FICON link incident reporting” on page 357.

### ***IOSHC112E***

#### **ANALYSIS OF COMMAND RESPONSE (CMR) TIME DETECTED ONE OR MORE CONTROL UNITS WITH AN EXCEPTION.**

IOS\_CMRTIME\_MONITOR IBM Health Check function, based on its configuration, has detected one or more control units with an inconsistent command response (CMR) time.

This new function uses the following configuration parameters to determine if an exception exists:

- ▶ THRESHOLD, value in milliseconds of the CU average CMR time.
- ▶ RATIO, value that determines the number of times is the average CMR time of a detected CU path greater than the lowest average CMR time for the same CU.

For example, having the following configuration: THRESHOLD(3); RATIO(5). CMR Time function will detect an exception if the following conditions are met:

- ▶ The average CMR time of the CU path is greater than 3 milliseconds.
- ▶ The average CMR time of the CU path is, at least, 5 times greater than the path with the lowest average CMR time for this control unit.

For more information, see 10.1.3, “IBM Health Checker: Command Response (CMR) Time Function” on page 351.

**Note:** Refer to message IOSHC113I to see the detail information about the exception detected with IOSHC112E. After you have the information about the affected control units and channel paths, use diagnostics tools to help determine the problem.

### ***IOSHC113I***

#### **COMMAND RESPONSE TIME REPORT**

This message is an IOSHC112E message extension. The message gives you detailed information about the CMR time exception detected by Health Checker function: IOS\_CMRTIME\_MONITOR.

The message has two main parts:

- ▶ The values of *Threshold* and *Ratio* from which the exception has been detected.
- ▶ Detailed information of the CU path detected.

For more information, see 10.1.3, “IBM Health Checker: Command Response (CMR) Time Function” on page 351.

*Example 10-4 IOSHC113I Command Response Time Report*

---

IOSHC113I COMMAND RESPONSE TIME REPORT

THE FOLLOWING CONTROL UNITS SHOW INCONSISTENT  
AVERAGE COMMAND RESPONSE (CMR) TIME BASED ON THESE

PARAMETERS:

THRESHOLD = threshold

RATIO = x

CMR TIME EXCEPTION DETECTED AT: time

CONTROL UNIT = cu

ND = node

|       | ENTRY | EXIT | CU   | I/O  | AVG |
|-------|-------|------|------|------|-----|
| CHPID | LINK  | LINK | INTF | RATE | CMR |
| chpid | e1    | d1   | intf | iort | cmr |

---


**Note:** IOSHC113I is a message extension of IOSHC112E. Refer to "IOSHC112E" on page 393 for more information.



# Part 5

# Appendixes





## **Example: Planning workflow**

This appendix shows an example of a workflow using the steps described in Chapter 4, “Planning the FICON environment” on page 77.

## Step 1: Documentation

Creating and maintaining documentation throughout all the planning, design, and implementation phases is important. Throughout this book we provide information regarding the various features and functions offered by System z servers and the FICON Directors (and under which conditions they are to be applied). That information can be used as a starting point for your documentation.

For documentation of our implementation steps for a cascaded FICON Director environment see Chapter 7, “Configuring a cascaded topology” on page 197, and Chapter 8, “Configuring FICON Directors” on page 245. Configuration worksheets for your FICON Director and FICON CTC environments can be found in Appendix B, “Configuration worksheets” on page 405.

## Step 2: Requirements

We based our requirements on the need to access non-business-critical data in a remote location from two z/OS LPARs in the main location. Isolating the new configuration from the existing one is essential to the solution. High availability is *not* a key requirement.

Based on our physical and logical inventory, these were the components involved in the solution:

- ▶ An IBM zEnterprise 196 in the main location:
  - Two z/OS V1R13 LPARs
  - Two FICON Express8S channels
- ▶ An IBM System Storage DS8000 in the remote location:
  - Two storage control units and two devices
  - Two 8 Gbps ports
- ▶ Two IBM System Storage SAN b-type family (FICON Directors):
  - SAN768B-2 with four 8 Gbps ports (in the main location)
  - SAN384B-2 with four 8 Gbps ports (in the remote location)

Figure A-1 displays the components used to create our new FICON environment.




Figure A-1 Results of the inventory effort

## Step 3: Context

Taking into consideration the existing and planned components (System z, FICON Director, and storage control device), we determined that the scenarios 4.4.2, “Moving to a high

bandwidth environment (FICON Express8/8S)” on page 83 and 4.4.3, “Migrating from a single site to a multi-site environment” on page 84, best fit our requirements.

In Figure A-2, ports B4 and C6 in SAN768B-2 and ports 00 and 20 in SAN384B-2 are used for ISLs. All ports run at data link rate of 8 Gbps.




Figure A-2 Our FICON components in the two locations

#### Step 4: Topologies and supported distances

We implemented a cascaded FICON Director topology (see 4.5.3, “Cascaded” on page 87 for more information). The two locations were less than the maximum distance of 10 km apart, therefore a WDM environment was not required. For more information about supported distances, refer to Table 4-1 on page 85.

The transceiver type required for the ISLs in a cascaded FICON Director topology is long wavelength (LX). LX transceivers are generally recommended for all connections, which we used for all our ports, as shown in Figure A-3.




Figure A-3 Our cascaded FICON Director topology

#### Step 5: Convergence

In our FICON environment, we did not set up an intermix environment. However, the FICON Directors do offer this possibility, if needed. Because we were using existing environments, other relevant factors for the FICON fabric such as power consumption, cooling, and space inside the data center were not a concern.

Our FICON environment was connected to an isolated network within a secured area, therefore we only used *username and password* authentication and authorization. However, we highly recommend that you change all the default passwords for all default accounts.

To isolate the new configuration from the existing configuration, we secured the fabric using zoning and the fabric configuration server with strict SCC policies.

### Step 6: Management

For management purposes, we used the IBM Network Advisor for the entire FICON environment including setup and monitoring, as illustrated in Figure A-4.

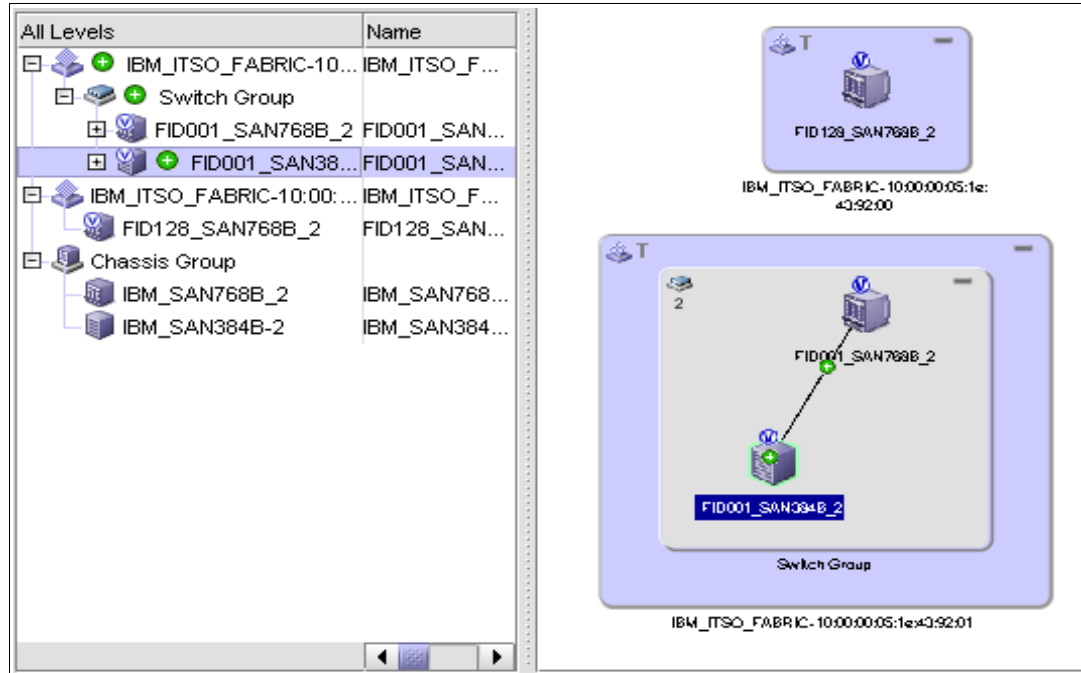


Figure A-4 FICON Management using IBM Network Advisor

Our FICON Director management environment is shown in Figure A-5. The figure also shows the IP addresses that we used for our FICON Directors, IBM Network Advisor server, and firewall/router.

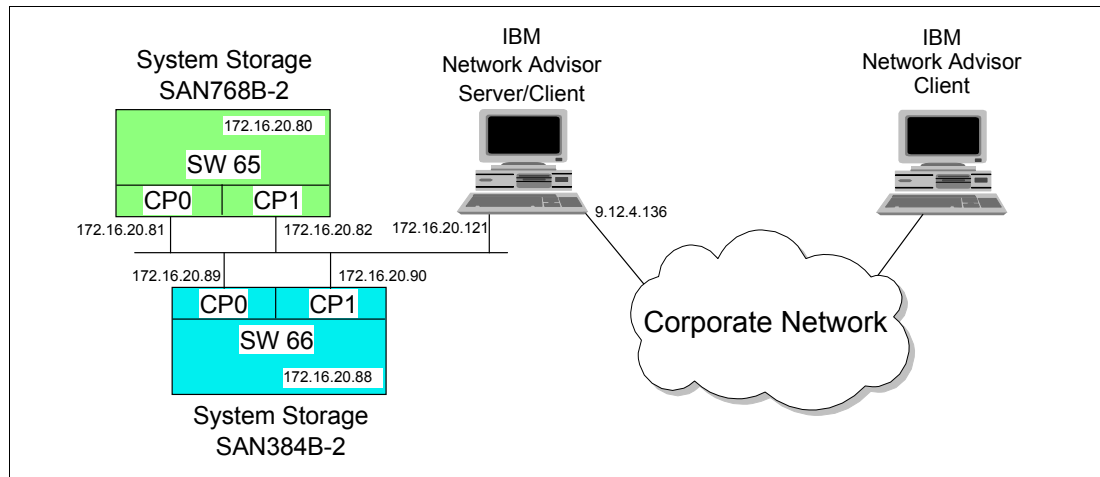


Figure A-5 Our FICON Director management environment



We also used the CUP and SA for I/O to manage and control System z FICON connectivity.

### Step 7: Virtualization

In System z, we had two z/OS LPARs. Although the two z/OS LPARs use the same CSS, each FICON channels will be defined to an additional CSS (spanned) for future use. The channels were on two different FICON features. We defined redundant paths from both z/OS LPARs to the storage control unit.




Figure A-6 Virtualization

For the DS8000, we set up two ports for redundant access to all devices. We had two logical control units inside the DS8000, and we defined the DS8000 channels to different FICON features.

Our FICON Directors were, by default, set up with virtualization support and we used this technology to define port addresses beyond the FICON port addressing range. We created a new logical switch to be able address the 48-port line cards. We also assigned our ports to two different cards for redundancy. The two virtual FICON Directors were interconnected with ISLs.

Figure A-7 on page 402 and Figure A-8 on page 402 show the worksheets that were used to configure the FICON connectivity for the z196, the two FICON Directors, and the DS8000. Figure A-7 on page 402 displays the configuration worksheet for FICON Director 1.

| FICON Director Configuration Worksheet  |             |              |                     |   |  |              |       |               |                      |
|---|-------------|--------------|---------------------|---|--|--------------|-------|---------------|----------------------|
| FICON Director Manufacturer: IBM        |             |              |                     | Type: 2499                                    |  | Model: 816   |       | S/N: 02-8D014 |                      |
| HCD Defined Switch ID x65 (Switch ID)   |             |              |                     | Cascaded Directors No Yes x                   |  |              |       |               |                      |
| FICON Director Domain ID x65 (Switch @) |             |              |                     | Corresponding Cascaded Director Domain ID x66 |  |              |       |               |                      |
|   |             |              |                     | Fabric Name Cascaded Fabric 1                 |  |              |       |               |                      |
| FICON Director Ports                    |             |              |                     |   | Attached N_Ports / E_Ports (CU, CPC, or ISL) |              |       |               |                      |
| Slot Number                             | Port Number | Port Address | Laser Type: LX / SX | Port Name                                     | Node Type CU/CHNL/FD                         | Machine Type | Model | Serial Number | ISL CU I/F CPC CHPID |
| 9                                       | 0           | 6592         | LX                  | z196 PCHID 399                                | CHNL   | 2817         | M32   | 02-B3BD5      | CPC CHPID            |
| 10                                      | 0           | 65A3         | LX                  | z196 PCHID 3E1                                | CHNL   | 2817         | M32   | 02-B3BD5      | CPC CHPID            |
| 11                                      | 0           | 65B4         | LX                  | ISL1 to SAN384B-2                             | FD   | 2499         | 416   | 02-2C00Z      | ISL                  |
| 1                                       | 26          | 65C6         | LX                  | ISL2 to SAN384B-2                             | FD   | 2499         | 416   | 02-2C00Z      | ISL                  |

Figure A-7 Configuration worksheet for FICON Director 1

Figure A-8 displays the configuration worksheet for FICON Director 2.

| FICON Director Configuration Worksheet  |             |              |                     |   |  |              |       |               |                      |
|---|-------------|--------------|---------------------|---|--|--------------|-------|---------------|----------------------|
| FICON Director Manufacturer: IBM        |             |              |                     | Type: 2499                                    |  | Model: 416   |       | S/N: 02-2C00Z |                      |
| HCD Defined Switch ID x66 (Switch ID)   |             |              |                     | Cascaded Directors No Yes x                   |  |              |       |               |                      |
| FICON Director Domain ID x66 (Switch @) |             |              |                     | Corresponding Cascaded Director Domain ID x65 |  |              |       |               |                      |
|   |             |              |                     | Fabric Name Cascaded Fabric 1                 |  |              |       |               |                      |
| FICON Director Ports                    |             |              |                     |   | Attached N_Ports / E_Ports (CU, CPC, or ISL) |              |       |               |                      |
| Slot Number                             | Port Number | Port Address | Laser Type: LX / SX | Port Name                                     | Node Type CU/CHNL/FD                         | Machine Type | Model | Serial Number | ISL CU I/F CPC CHPID |
| 7                                       | 1           | 6641         | LX                  | DS8K Port 0133                                | CU   | 2107         | 932   | 75-L3000      | CU I/F               |
| 1                                       | 17          | 6611         | LX                  | DS8K Port 0333                                | CU   | 2107         | 932   | 75-L3000      | CU I/F               |
| 1                                       | 0           | 6600         | LX                  | ISL1 to SAN768B-2                             | FD   | 2499         | 816   | 02-8D014      | ISL                  |
| 2                                       | 0           | 6620         | LX                  | ISL2 to SAN768B-2                             | FD   | 2499         | 816   | 02-8D014      | ISL                  |

Figure A-8 Configuration worksheet for the FICON Director 2

For high integrity, we set up Insistent Domain IDs and strict SCC policy. We used one port-based zone for all our FICON channels.

### Step 8: Performance

The short distance between the two locations allowed us to keep the default buffer credit assignment in the FICON Directors. For high performance, we used the zHPF feature in both the System z server and the DS8000. For better performance of the fabric links, we used lossless DLS with frame-based trunking.

We also used port fencing to ensure that no loss of performance occurs due to excessive errors on the interfaces, and to make certain that the fabric remains stable.

To monitor and evaluate the utilization of the paths and the traffic patterns, we used RMF on our System z server.

### Step 9: Prerequisites and interoperability

We ensured that all System z, DS8000, and FICON Director prerequisites and interoperability requirements are fulfilled, for example:

- ▶ zEnterprise 196:
  - Feature code for FICON Express8 LX FC3325
  - Feature code for FICON Express8S LX FC0409 (for use with Seneca I/O drawer FC4003)
- ▶ DS8000:
  - Feature code for zHPF FC0709 and FC7092
  - Feature code for FICON attachment (FC0703 and FC7091)
- ▶ SAN768B-2 and SAN384B-2:
  - Feature code for the FICON CUP management FC7886

See references in 4.10, “Prerequisites and interoperability” on page 112.

### Step 10: Physical connectivity

Our set up was based on LX FICON transceivers with standard LC duplex connectors, therefore we used 9 μm single-mode fiber optic cables (OS2 fiber type), as shown in Figure A-9.




Figure A-9 Physical connectivity

For the cabling infrastructure we used IBM Fiber transport service products with single mode fiber.

### Conclusion

After all these steps were completed, our final FICON design appeared as shown in Figure A-10 on page 404.




Figure A-10 Our final FICON design



# B

## Configuration worksheets

This appendix contains the following worksheets:

- ▶ FICON Director Configuration Worksheet

Use this worksheet to document the layout of your FICON Director. It can be applied as a tool to help you understand how the ports are allocated for configuration and problem determination purposes.

- ▶ FICON CTC Image-ID Worksheet

Use this worksheet as an aid in planning and documenting the association between LPARs, MIF IDs, Logical Channel subsystems, and CTC Image-IDs.











# Configuration and definition tools

We provide an overview of tools that you can use to define paths between a System z server, FICON directors, and a DASD control unit and a device. The tools are:

- ▶ Hardware Configuration Definition
- ▶ Hardware Configuration Manager
- ▶ CHPID Mapping Tool
- ▶ I/O Configuration Program
- ▶ Stand-Alone I/O Configuration Program
- ▶ IOCP statements and keywords used for FICON

HCD is a program that runs in a z/OS image. HCM and CMT are PC-based programs. IOCP is a stand-alone program running in a partition on a System z server.

HCD, HCM, CMT, and IOCP are briefly described in this appendix. Also included are descriptions of statements and keywords that are used to define a FICON point-to-point, switched or cascaded topology.

Unlike other platforms, with System z all paths to an I/O device must be defined (there is no auto-detection or auto-configuration).

**Note:** However there is now an option in HCD called zDAC (Option 1.6 - Discovered new and changed control units and I/O devices) which allows you to discover undefined FICON storage devices (DASD and tape) connected to the processor using a FICON director.

The characteristics of the channel paths, the control units, and the devices have to be specified and saved in an I/O Configuration Data Set (IOCDS), which is used by a System z server to load the configuration into its Hardware System Area (HSA). The IOCDS file is stored on the Support Element. Although HCD and HCM can create an IOCDS for each physical CPC present in a sysplex environment, IOCP can create an IOCDS only for the physical system it is running on.

**Note:** However using HCD (Option 2.11 - Build and manage s/390 microprocessor IOCDs and IPL attributes), having all of the applicable CPCs defined in the IODF and having the Network name and CPC name defined in the IODF, which matches the SNA address of the CPC (USIBMSC.SCZP301), an IOCP deck can be built and written to any of these CPCs IOCDs.

The IOCDs contains statements and keywords that specify the characteristics of a channel, control unit, and device. These statements and keywords are described in “IOCP statements and keywords used for FICON” on page 412.

Refer to *IBM zEnterprise 196 Configuration Setup*, SG24-7834, for additional information about how to set up the System z server.

## Hardware Configuration Definition

The Hardware Configuration Definition (HCD) supplies an interactive dialog to generate your I/O definition file (IODF) and subsequently your Input/Output Configuration Data Set (IOCDs). We strongly recommended that you use HCD to generate your IODF, as opposed to writing your own IOCP statements. The validation checking that HCD performs as you enter data helps to eliminate errors before you implement your I/O configuration.

Refer to *HCD User's Guide*, SC33-7988, for detailed information about HCD.

## Hardware Configuration Manager

The Hardware Configuration Manager (HCM) is a PC-based client/server interface to HCD that combines the logical and physical aspects of hardware configuration management. In addition to the logical connections, you can also manage the physical aspects of a configuration. For example, you can effectively manage the flexibility offered by the FICON infrastructure.

All updates to your configuration can be done through HCM's intuitive graphical user interface and, most important, due to the client/server relationship with HCD, all changes to the logical I/O configuration are written into the IODF and fully validated and checked for accuracy and completeness by HCD.

If you are licensed to use HCM, the install file that runs on the workstation can be found in dataset SYS#.SEEQINST. The member name is EEQINSTM and is downloaded to the workstation and installed.

Two additional started tasks must also run on the host:

- ▶ CBDQDISP (known as the HCM-HCD DISPatcher job for TCP/IP connection)
- ▶ CBDQAJSK (known as the HCM-HCD Agent Jcl SKeleton job for each individual host session between HCM and HCD)

Refer to *HCM User's Guide*, SC33-7989, for detailed information about HCM.

## CHPID Mapping Tool

The CHPID Mapping Tool (CMT) provides a mechanism for customizing the CHPID assignments for System z servers to avoid connecting critical paths to single points of failure. Additional enhancements have been built into the CMT to help with the requirements of z196

and z114 servers, thereby providing a mechanism to map CHPIDS to PCHIDS, and the best availability recommendations for installed features and defined configurations.

More details and informations about the CHPID Mapping Tool are available on IBM Resource Link® at the following URL:

<https://www.ibm.com/servers/resourceLink>

After you have registered, the CMT program and the CMT user's guide can be downloaded to your workstation.

Simply click your machine type to access the available resources for that System z server. The latest version of CHPID Mapping Tool download link can be found under resources.

## I/O Configuration Program

The I/O Configuration Program (ICP IOCP) is a program that defines to a CPC all of the available I/O devices and the channel paths.

In the CPC, the input/output (I/O) subsystem, which controls channel operations, requires specific data about the hardware I/O configuration. To define the I/O configuration for the I/O subsystem, Input/Output Configuration Program (IOCP) needs to be used.

- ▶ To run IOCP, you need to specify:
- ▶ Logical partitions
- ▶ Channel paths on the central processor complex (CPC) and their assignment to logical partitions
- ▶ Control units attached to the channel paths
- ▶ I/O devices assigned to the control units

Also, to meet changing I/O requirements, you can replace an existing I/O configuration with a new configuration by running IOCP. All control units and I/O devices that attach to the CPC must be defined to the channel subsystem by using IOCP.

**Note:** IOCP is generally invoked using HCD (Option 2 - Activate or process configuration data) or through running JCL, but we strongly recommend that you use HCD to generate your IODF as opposed to writing your own IOCP statements.

## Stand-Alone I/O Configuration Program

Sometimes it is necessary to create a Stand-Alone IOCP file that contains the IOCP statements built from a production IODF and required to perform a Power-on Reset of a CPC.

This situation can occur if you were installing a new CPC in a site that has no connectivity to the system where you built the production IODF.

The processor requires:

- ▶ Build the production IODF (HCD Option 2.1 - Build production I/O definition file)
- ▶ Create the IOCP input statements (HCD Option 2.3 - Build IOCP input dataset)
- ▶ Transfer the IOCP input statements to a USB flash drive
- ▶ Use these statements as input to the stand-alone IOCP process on the Hardware Management Console (HMC) for the new CPC

Refer to *IOCP User's Guide*, SB10-7037, for detailed information about IOCP.

## IOCP statements and keywords used for FICON

This section describes the required keyword definitions and rules for a FICON configuration.

Refer to *IOCP User's Guide*, SB10-7037, for a detailed description of statements and keywords. The latest version of this publication is available on IBM Resource Link under Library → Publications

<https://www.ibm.com/servers/resourceLink>

### IOCP and z/OS I/O configuration

The Hardware Configuration Definition (HCD) defines the I/O configuration to the z/OS operating system and the I/O subsystem. You do not need to run IOCP directly to define the I/O configuration to the I/O subsystem. However, HCD does run IOCP when it needs to write an IOCDS:

|                 |  |
|-----------------|--|
| <b>ID</b>       | Statement (optional) that defines the heading data for configuration reports. Only one ID statement is allowed.                  |
| <b>RESOURCE</b> | Statement (required) that defines the logical partitions and logical channel subsystems. Only one RESOURCE statement is allowed. |
| <b>CHPID</b>    | Statements (optional) that define channel paths.   |
| <b>CNTLUNIT</b> | Statements (optional) that define control units. Each control unit must be defined in a separate CNTLUNIT statement.             |
| <b>IODEVICE</b> | Statements (optional) that define I/O devices. As many as 256 I/O devices can be defined in an IODEVICE statement.               |

### ID statement

The ID statement is an optional statement that can specify the following information:

- ▶ Identification information printed in the heading of IOCP configuration reports on the ID1 and ID2 lines.
- ▶ Data set name for each IOCDS that appears on the Input/Output Configuration window of the Support Element for a CPC
- ▶ Set of machine limits and rules to be enforced by IOCP when processing the I/O configuration
- ▶ Name of the local system.

In addition to other keywords allowed in the ID statement, only keywords used later in this chapter are described here:

### **SYSTEM**

The SYSTEM keyword specifies the machine limits and rules that are used for verification of input data set. The device type of the server using the IOCDS has to be specified in the ID keyword. The TOK keyword provides information about the source file used to build the IOCDS.

### **RESOURCE statement**

The RESOURCE is a required statement that defines:

- ▶ Logical partition names

- ▶ Multiple Image Facility (MIF) image ID numbers assigned to the logical partitions
- ▶ Logical channel subsystems (CSSs)
- ▶ The maximum number of devices that can be dynamically defined for a dynamic-capable IOCDS
- ▶ Reserved logical partitions, for a dynamic-capable IOCDS.

The RESOURCE statement must:

- ▶ Appear only once in an IOCP input file
- ▶ Follow the ID statement (if specified)
- ▶ Precede all CHPID, CNTLUNIT, and IODEVICE statements
- ▶ Include every logical partition in the configuration
- ▶ Have either the PART or PARTITION parameter specified.

### **PART**

The PART keyword allows you to specify an LPAR name and to assign the MIF ID to a partition. Optionally, a CSS number can be assigned to the logical partition.

### **CHPID statement**

The CHPID is an optional statement that defines:

- ▶ Channel path identifiers (*chpids*)
- ▶ Physical channel identifiers (*pchids*)
- ▶ Channel path types
- ▶ *chpid* assignment to ESCON and FICON Directors
- ▶ Connecting *chpids*
- ▶ Channel path parameter
- ▶ I/O cluster name owning a *chpid*
- ▶ Adapter identifiers (AIDs), ports, and connecting system names
- ▶ Availability of *chpids* to logical partitions
- ▶ Reconfigurability of *chpids* among logical partitions
- ▶ Shared *chpids* for logical partitions
- ▶ Spanned *chpids* for logical channel subsystems.

When specified, the CHPID statement must:

- ▶ Follow the ID statement (if specified) and the RESOURCE statement
- ▶ Precede any CNTLUNIT statements that refer to the *chpid* it defines
- ▶ Collectively, specify all *chpids* to use in the configuration.

Note that not all keywords allowed in the CHPID statement are described here; only keywords used to define a FICON channel (FC) are listed:

|               |   |
|---------------|---|
| <b>PATH</b>   | The PATH parameter allows us to specify the CSSs and the CHPID number that will be assigned to a channel path. The PCHID keyword will be used to identify the location of the physical channel in the processor.  |
| <b>TYPE</b>   | The CHPID type for FICON channels in a point-to-point configuration is required and must be specified as TYPE=FC. This indicates that the FICON channel will operate in FICON native (FC) mode.   |
| <b>SWITCH</b> | The Switch keyword is required, and it specifies an arbitrary number for the FICON Director to which the channel path is assigned. Note that the switch number defined in the IOCDS is not the same as the switch address that is set up in the FICON Director. The switch address is the Domain ID specified in the FICON Director. It is recommended that you have the same value for the switch number and the switch address. |

**Note:** The switch number and switch address are defined as a hexadecimal value in IODF. The Domain ID defined in the FICON Director is a decimal value. For example, a switch address of 65 in IODF equals a Domain ID of 101.

### **CNTLUNIT statement**

The CNTLUNIT is an optional statement that describes:

- ▶ The characteristics of the control unit
- ▶ The channel paths to which the control unit is attached
- ▶ The unit addresses the control unit recognizes.

Note that not all of the keywords allowed in the CNTLUNIT statement are described here; only keywords used to define a control unit connected to a FICON channel (FC) are listed.

**PATH** The PATH keyword specifies the channel path in each channel subsystem (CSS) attached to the control unit.

**LINK** The LINK keyword is required to define the destination port address where the control unit is connected to the FICON Director. The link address can be specified with one or two bytes. Specifying a two-byte link address requires a special feature (High Integrity Fabric) and setup in the FICON Director.

**CUADD** For FICON native (FC) channel paths, the logical address (or Logical Control Unit - LCU) is specified as two hexadecimal digits in the range 00–FE.

Not all FICON control units support logical addressing. To determine a product's logical addressing information, contact your System Service Representative supporting the FICON control units.

**UNITADD** The UNITADD keyword specifies the unit addresses of the I/O devices that the control unit recognizes.

### **IODEVICE statement**

The IODEVICE is an optional statement that describes the following:

- ▶ The I/O device number
- ▶ The device characteristics
- ▶ The control units to which the device is assigned.

#### ***CUNUMBR***

The CUNUMBR keyword specifies the control unit number to which the device is attached.



## Configuring the DS8000 for FICON

The DS8000 supports FICON attachment to FICON channels on zEnterprise 196 (z196), zEnterprise 114 (z114), and System z10 (z10 EC, z10 BC). The following tasks must be done when configuring a DS8000 for use in a FICON environment.

- ▶ Activate the license keys for the storage unit.
- ▶ Configure the installed disk drives in RAID arrays and in Count-Key-Data (CKD) format.
- ▶ Create System z logical control units (LCUs), CKD base volumes, and Parallel Access Volume (PAV) aliases.
- ▶ Configure I/O ports for FICON use.

The actual configuration can be done using either the DS8000 Storage Manager Graphical User Interface (GUI) or the DS8000 Command-Line Interface (CLI), or a mixture of both. A novice user might prefer to use the GUI, and a more experienced user might use the CLI, particularly for the more repetitive tasks, such as creating large numbers of volumes.

In this appendix, we describe in detail how to configure the DS8000 I/O ports to communicate with a System z server using the FICON protocol. For information about how to perform the other tasks, see *IBM System Storage DS8000 Architecture and Implementation*, SG24-8886.

## DS8000 licenses for FICON

Support for FICON attachment is an optional feature of the DS8000 and is available with the FICON attachment licensed feature indicator feature numbers 7091 and 0703 and corresponding DS8000 series function authorization (2396-LFA FICON attachment feature number 7091). You must order and activate the FICON attachment feature to use your DS8000 in a FICON environment.

The following DS8000 features are optional but highly recommended for FICON environments:

- ▶ High performance FICON for System z (zHPF)

Support for High Performance FICON is an optional feature of DS8800 model 951 and DS8700 model 941 and is available with the High Performance FICON licensed feature indicator feature numbers 7092 and 0709 and corresponding DS8000 series function authorization (2396-LFA High Performance FICON feature number 7092). zHPF also requires the purchase of the FICON attachment licensed feature.

- ▶ Parallel Access Volumes (PAV)

PAV is an optional feature on the DS8000 series and is available with the PAV indicator feature numbers 78xx and 0780 and corresponding DS8000 series function authorization (2396-LFA PAV feature numbers 78xx). PAV also requires the purchase of the FICON attachment licensed feature.

- ▶ HyperPAV

HyperPAV is an optional feature on the DS8000 series and is available with the HyperPAV indicator feature numbers 7899 and 0782 and corresponding DS8000 series function authorization (2396-LFA HyperPAV feature number 7899). HyperPAV also requires the purchase of PAV licensed features and the FICON attachment licensed feature.

- ▶ DS8000 I/O Priority Manager

The DS8000 now features I/O Priority Manager, which enables more effective storage consolidation and performance management with the ability to align quality of service levels to separate workloads in the system. This capability is exclusive to the DS8000.

The DS8000 I/O Priority Manager feature is available for open systems and System z. Support for DS8000 I/O Priority Manager is an optional feature for DS8800 model 951 and DS8700 model 941, and is available with the I/O priority Manager licensed feature indicator DS8000 Function Authorization (239x-LFA) I/O Priority Manager feature number 784x and 0784, and corresponding DS8000 Function Authorization (2396-LFA) I/O Priority Manager feature number 784x.

- ▶ Easy Tier

IBM System Storage Easy Tier is designed to help automate data placement throughout the DS8800 disk pool (including multiple drive tiers) to intelligently align the system with current workload requirements. This includes the ability for the system to automatically and non disruptively relocate data across drive tiers, and the ability to manually relocate full volumes. The potential benefit is to align performance of the system with the appropriate application workloads.

Support for IBM System Storage Easy Tier is an optional feature of the DS8800 model 951 and DS8700 model 941, and is available with the IBM System Storage Easy Tier licensed feature indicator feature numbers 7083 and 0713 and corresponding DS8800 Function Authorization (2396-LFA) IBM System Storage Easy Tier feature number 7083.



# Configuring DS8000 I/O ports

DS8000 Fibre Channel/FICON host adapters (HA) support both Fibre Channel Protocol (FCP) for open systems attachment and FICON protocol for System z attachment, but not simultaneously. You can configure each port independently to run either FCP or FICON protocol. The port configuration can be changed as required. After configured, the ports can be connected to either System z FICON ports or to FICON director ports. When a port is configured for FICON, it also supports zHPF if the required zHPF licensed feature is installed on the DS8000.

The DS8700 offers 4-port 4 Gbps and 8 Gbps Fibre Channel/FICON host adapters, available both in longwave and shortwave. The 4 Gbps ports independently auto-negotiate to either 4, 2, or 1 Gbps link speed. The 8 Gbps ports independently auto-negotiate to either 8, 4, or 2 Gbps link speed.

The DS8800 offers 4-port and 8-port 8 Gbps Fibre Channel/FICON host adapters, available both in longwave and shortwave. Each port independently auto-negotiates to either 8, 4, or 2 Gbps link speed.

## I/O port layout

A DS8000 can have I/O adapters installed in two racks: rack 1 (R1) and rack 2 (R2). Rack 1 is the base rack with the storage controllers and the HMC in it. Rack 2 is the first expansion rack.

Both R1 and R2 can have four I/O enclosures (I1-I4 or B1-B4), each of which consists of six card slots (C1-C6). Host adapter (HA) cards, which are used for host access, are installed in card slots C1, C2, C4, and C5. Card slots C3 and C6 contain device adapter (DA) cards, which are used for connecting the backend disks to the DS8000 controllers.

Each HA card on a DS8700 and previous DS8000 models has four ports (T0-T3). The HA cards on a DS8800 can have either four or eight ports (T0-T7).

DS8000 I/O ports are identified by *location codes* and *port numbers*. The location codes are of format Rx-lx-Cx-Tx or Rx-Bx-Cx-Tx. They identify the physical location of the hardware components and are mainly used by hardware service representatives. The I/O port numbers are of format lnnnn. (They are sometimes written without the l in front.) An *interface ID* or System Adapter Identifier (SAID) is also associated with each port. Normally it is the same as the port number.<sup>1</sup> Port numbers and interface IDs are used for system configuration, management, and reporting purposes.

We next describe the port numbering for the DS8700 and DS8800 storage systems in more detail.

---

<sup>1</sup> There are some exceptions in certain remote copy configurations. See IBM Technical Document TD105862.

## DS8700 port layout

Figure D-1 shows the location codes of HA ports for the DS8700 and earlier DS8000 models. The numbering is identical for both rack 1 and rack 2.

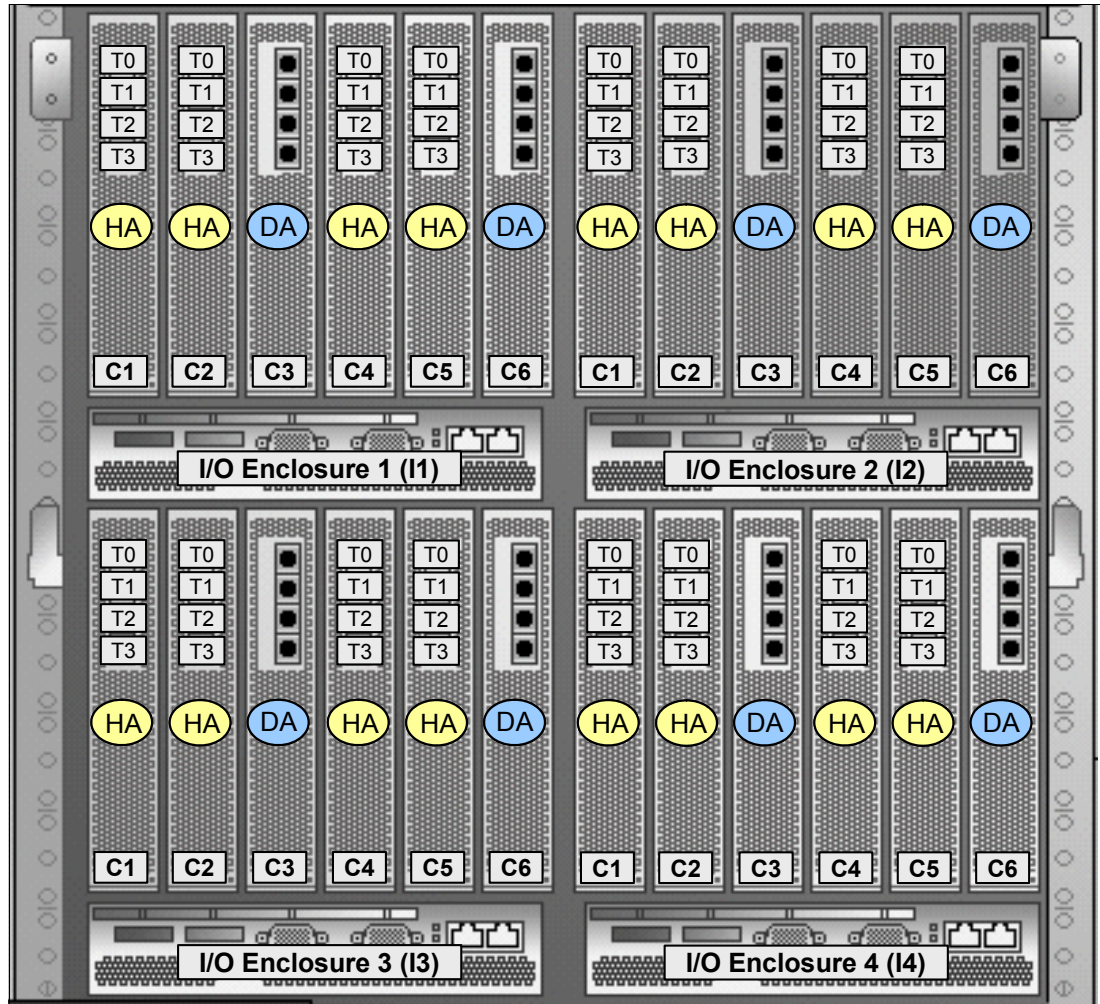


Figure D-1 Port location codes for DS8700 (rear view)

Table D-1 on page 419 and Table D-2 on page 419 show the HA port numbers for rack 1 and rack 2, respectively. The port numbers are written here without the I in front.

Table D-1 Port numbers for DS8700 rack1 (rear view)

| I1   |      |    |      |      |    | I2   |      |    |      |      |    |
|------|------|----|------|------|----|------|------|----|------|------|----|
| 0000 | 0010 |    | 0030 | 0040 |    | 0100 | 0110 |    | 0130 | 0140 |    |
| 0001 | 0011 | D  | 0031 | 0041 | D  | 0101 | 0111 | D  | 0131 | 0141 | D  |
| 0002 | 0012 | A  | 0032 | 0042 | A  | 0102 | 0112 | A  | 0132 | 0142 | A  |
| 0003 | 0013 |    | 0033 | 0043 |    | 0103 | 0113 |    | 0133 | 0143 |    |
| C1   | C2   | C3 | C4   | C5   | C6 | C1   | C2   | C3 | C4   | C5   | C6 |
| 0200 | 0210 |    | 0230 | 0240 |    | 0300 | 0310 |    | 0330 | 0340 |    |
| 0201 | 0211 | D  | 0231 | 0241 | D  | 0301 | 0311 | D  | 0331 | 0341 | D  |
| 0202 | 0212 | A  | 0232 | 0242 | A  | 0302 | 0312 | A  | 0332 | 0342 | A  |
| 0203 | 0213 |    | 0233 | 0243 |    | 0303 | 0313 |    | 0333 | 0342 |    |
| C1   | C2   | C3 | C4   | C5   | C6 | C1   | C2   | C3 | C4   | C5   | C6 |
| I3   |      |    |      |      |    | I4   |      |    |      |      |    |

Table D-2 Port numbers for DS8700 rack 2 (rear view)

| I1   |      |    |      |      |    | I2   |      |    |      |      |    |
|------|------|----|------|------|----|------|------|----|------|------|----|
| 0400 | 0410 |    | 0430 | 0440 |    | 0500 | 0510 |    | 0530 | 0540 |    |
| 0401 | 0411 | D  | 0431 | 0441 | D  | 0501 | 0511 | D  | 0531 | 0541 | D  |
| 0402 | 0412 | A  | 0432 | 0442 | A  | 0502 | 0512 | A  | 0532 | 0542 | A  |
| 0403 | 0413 |    | 0433 | 0443 |    | 0503 | 0513 |    | 0533 | 0543 |    |
| C1   | C2   | C3 | C4   | C5   | C6 | C1   | C2   | C3 | C4   | C5   | C6 |
| 0600 | 0610 |    | 0630 | 0640 |    | 0700 | 0710 |    | 0730 | 0740 |    |
| 0601 | 0611 | D  | 0631 | 0641 | D  | 0701 | 0711 | D  | 0731 | 0741 | D  |
| 0602 | 0612 | A  | 0632 | 0642 | A  | 0702 | 0712 | A  | 0732 | 0742 | A  |
| 0603 | 0613 |    | 0633 | 0643 |    | 0703 | 0713 |    | 0733 | 0742 |    |
| C1   | C2   | C3 | C4   | C5   | C6 | C1   | C2   | C3 | C4   | C5   | C6 |
| I3   |      |    |      |      |    | I4   |      |    |      |      |    |

## DS8800 port layout

Figure D-2 shows the location codes of HA ports for the DS8800. The numbering is identical for rack 1 and rack 2. Note that card slots C2 and C5 are not used on the DS8800.

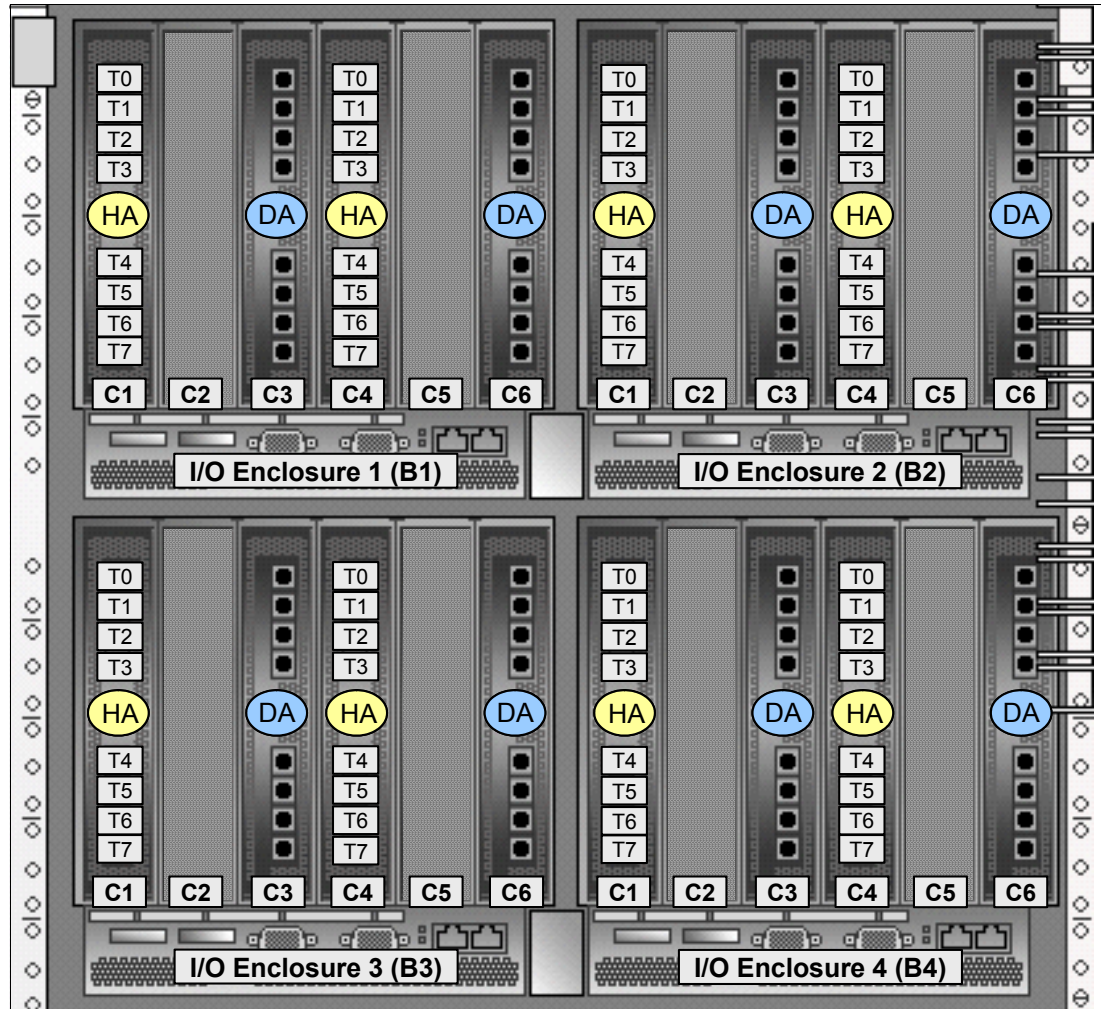


Figure D-2 Port location codes for DS8800 (rear view)

Table D-3 on page 421 and Table D-4 on page 421 show the HA port numbers for rack 1 and rack 2, respectively. The port numbers are written here without the I in front.

Table D-3 Port numbers for DS8800 rack1 (rear view)

| B1   |    |    |      |    |    | B2   |    |    |      |    |    |
|------|----|----|------|----|----|------|----|----|------|----|----|
| 0000 |    |    | 0030 |    |    | 0100 |    |    | 0130 |    |    |
| 0001 |    |    | 0031 |    |    | 0101 |    |    | 0131 |    |    |
| 0002 |    |    | 0032 |    |    | 0102 |    |    | 0132 |    |    |
| 0003 |    | D  | 0033 |    | D  | 0103 |    | D  | 0133 |    | D  |
| 0004 |    | A  | 0034 |    | A  | 0104 |    | A  | 0134 |    | A  |
| 0005 |    |    | 0035 |    |    | 0105 |    |    | 0135 |    |    |
| 0006 |    |    | 0036 |    |    | 0106 |    |    | 0136 |    |    |
| 0007 |    |    | 0037 |    |    | 0107 |    |    | 0137 |    |    |
| C1   | C2 | C3 | C4   | C5 | C6 | C1   | C2 | C3 | C4   | C5 | C6 |
| 0200 |    |    | 0230 |    |    | 0300 |    |    | 0330 |    |    |
| 0201 |    |    | 0231 |    |    | 0301 |    |    | 0331 |    |    |
| 0202 |    |    | 0232 |    |    | 0302 |    |    | 0332 |    |    |
| 0203 |    | D  | 0233 |    | D  | 0303 |    | D  | 0333 |    | D  |
| 0204 |    | A  | 0234 |    | A  | 0304 |    | A  | 0334 |    | A  |
| 0205 |    |    | 0235 |    |    | 0305 |    |    | 0335 |    |    |
| 0206 |    |    | 0236 |    |    | 0306 |    |    | 0336 |    |    |
| 0207 |    |    | 0237 |    |    | 0307 |    |    | 0337 |    |    |
| C1   | C2 | C3 | C4   | C5 | C6 | C1   | C2 | C3 | C4   | C5 | C6 |
| B3   |    |    |      |    |    | B4   |    |    |      |    |    |

Table D-4 Port numbers for DS8800 rack 2 (rear view)

| B1   |    |    |      |    |    | B2   |    |    |      |    |    |
|------|----|----|------|----|----|------|----|----|------|----|----|
| 0400 |    |    | 0430 |    |    | 0500 |    |    | 0530 |    |    |
| 0401 |    |    | 0431 |    |    | 0501 |    |    | 0531 |    |    |
| 0402 |    |    | 0432 |    |    | 0502 |    |    | 0532 |    |    |
| 0403 |    | D  | 0433 |    | D  | 0503 |    | D  | 0533 |    | D  |
| 0404 |    | A  | 0434 |    | A  | 0504 |    | A  | 0534 |    | A  |
| 0405 |    |    | 0435 |    |    | 0505 |    |    | 0535 |    |    |
| 0406 |    |    | 0436 |    |    | 0506 |    |    | 0536 |    |    |
| 0407 |    |    | 0437 |    |    | 0507 |    |    | 0537 |    |    |
| C1   | C2 | C3 | C4   | C5 | C6 | C1   | C2 | C3 | C4   | C5 | C6 |
| 0600 |    |    | 0630 |    |    | 0700 |    |    | 0730 |    |    |
| 0601 |    |    | 0631 |    |    | 0701 |    |    | 0731 |    |    |
| 0602 |    |    | 0632 |    |    | 0702 |    |    | 0732 |    |    |
| 0603 |    | D  | 0633 |    | D  | 0703 |    | D  | 0733 |    | D  |
| 0604 |    | A  | 0634 |    | A  | 0704 |    | A  | 0734 |    | A  |
| 0605 |    |    | 0635 |    |    | 0705 |    |    | 0735 |    |    |
| 0606 |    |    | 0636 |    |    | 0706 |    |    | 0736 |    |    |
| 0607 |    |    | 0637 |    |    | 0707 |    |    | 0737 |    |    |
| C1   | C2 | C3 | C4   | C5 | C6 | C1   | C2 | C3 | C4   | C5 | C6 |
| B3   |    |    |      |    |    | B4   |    |    |      |    |    |

## Configuring I/O ports using the DS8000 GUI

The following steps describe how to set up DS8000 I/O ports for FICON using the DS8000 Storage Manager Graphical User Interface (DS GUI). The ports can also be configured using the DS8000 Command-Line Interface.

The configuration steps apply equally to point-to-point, switched, and cascaded FICON topologies. zHPF does not require any additional configuration.

The DS GUI server code resides on the DS8000 Hardware Management Console (HMC). On a DS8000 with version 6.1 or earlier code level, the DS GUI can be accessed from the Tivoli Storage Productivity Center (TPC) application by launching an Element Manager for the DS8000 in the TPC GUI. You can use the TPC which comes preinstalled on the System Storage Productivity Center (SSPC) workstation which is shipped with the DS8000, or you can use a TPC on another workstation that has network connectivity to the DS8000 HMC.

With version 6.2 and later DS8000 code levels, the DS GUI can also be launched directly from any workstation that has network connectivity to the DS8000 HMC using a supported web browser. Point the browser to `http://<IP_address_of_DS8000_HMC>:8451/DS8000` or to `https://<IP_address_of_DS8000_HMC>:8452/DS8000`. When the logon panel is displayed, type in the DS8000 management user ID and password.

After connected to the DS GUI, perform the following steps to configure the I/O ports.

**Note:** The screen images shown here were captured on a DS8800 with version 6.2 code. The corresponding panels on version 5 and earlier code levels look considerably different.

1. On the Overview panel, hover over the Home icon, and select **System Status**, as shown in Figure D-3. The System Status window opens.

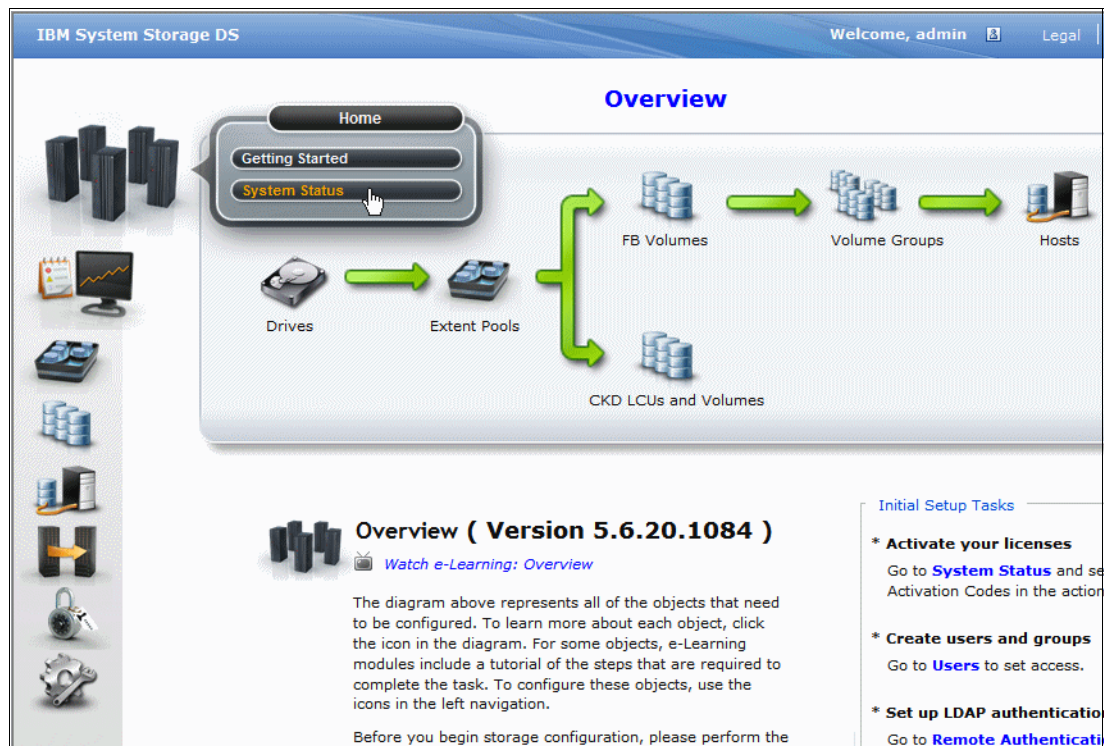


Figure D-3 DS8000 GUI Overview panel

- On the System Status window, Figure D-4, select the storage image for which you want to configure the ports by clicking the corresponding row. The selected row is highlighted in yellow.
- From the Action drop-down menu, select **Storage Image** → **Configure I/O ports**. The Configure I/O Ports window opens.

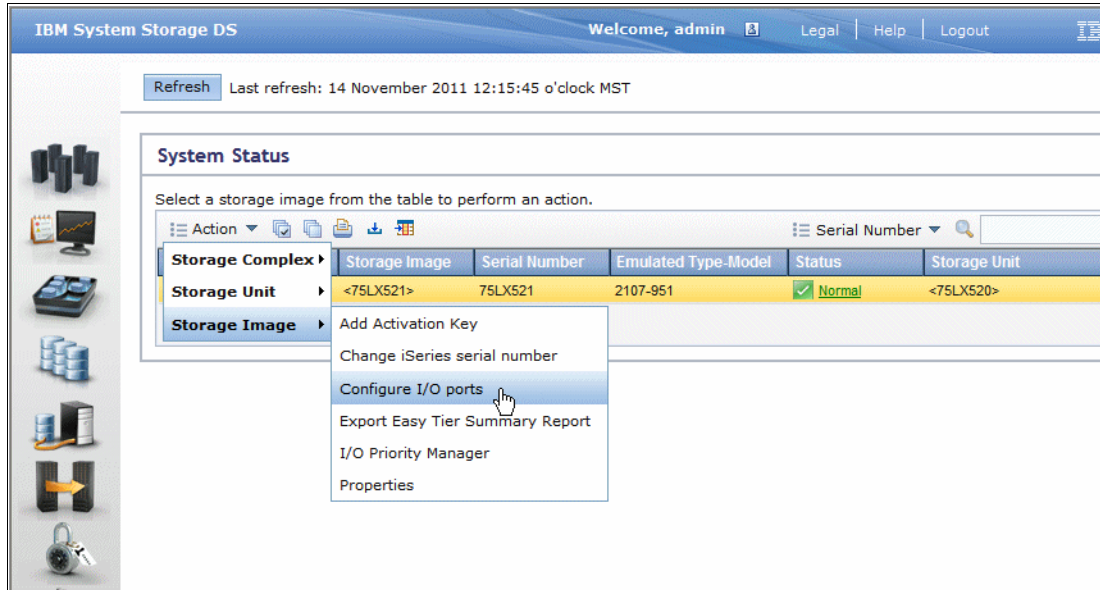


Figure D-4 DS8000 GUI System Status window

- On the Configure I/O Ports window, Figure D-5 on page 424, both the port number and location code of I/O ports are shown. Select the port you want to configure by clicking the corresponding row. Multiple port selection is supported. The selected rows are highlighted in yellow.
- Open the Action drop-down menu, and click **Change interface to FICON** to configure the selected ports to use FICON protocol.

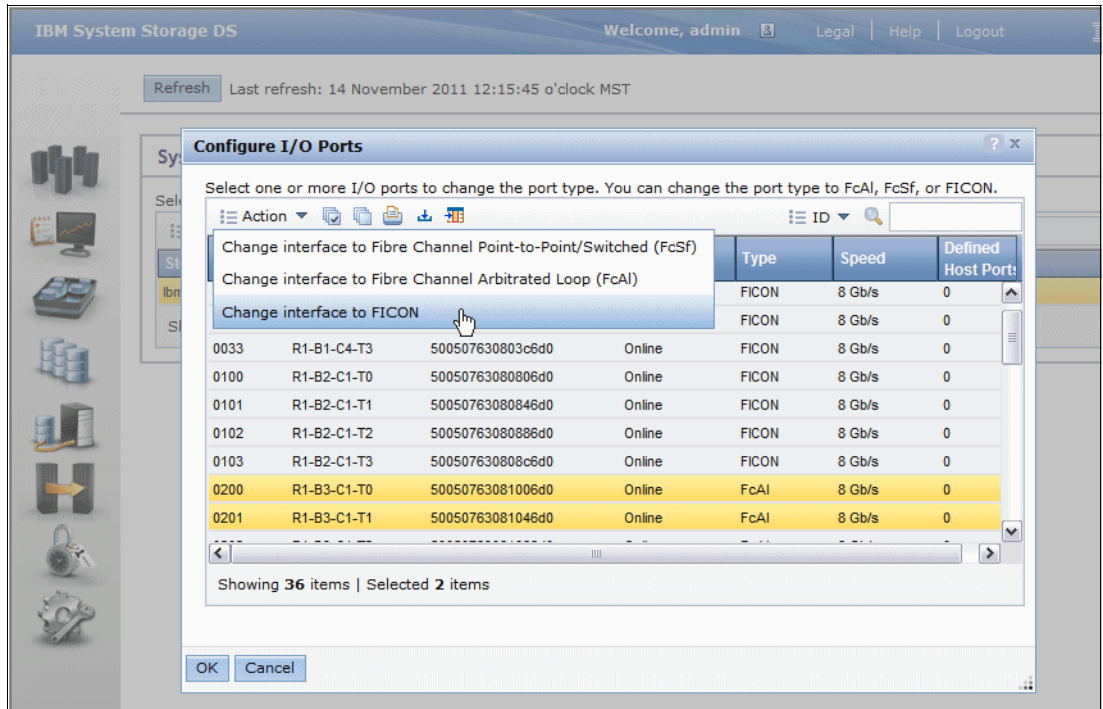


Figure D-5 DS8000 GUI Configure I/O Ports window

You get a warning message that the ports might become unusable by the hosts that are currently connected to them.

6. You can repeat step 3 to configure additional ports. When done, click **OK**.





# Using HMC and SE for problem determination information

This appendix explains how you can use the HMC and SE panels to obtain information when performing problem determination or verifying the FICON environment.

The following topics are discussed:

- ▶ E.1 HMC and SE information on page 426:
  - HMC and SE LIC version
  - HMC and SE default Userid and Password
  - How to logon HMC and SE in Tree Style User Interface
- ▶ E.2 Displaying CPC resource information on page 429:
  - Current IOCDS, IODF, and Activation Profile
  - Machine Type and Model
  - Frame Layout information
  - Channel Assignment with PCHID,CHPID,CSS, and FICON slots
- ▶ E.3 Displaying individual FICON channel information on page 435:
  - FICON channel status
  - Subchannel information
  - Control Unit definition information
  - Device and Link status
  - FICON LESB (Link Error Statistics Block) information
  - FICON channel Optical Power level
- ▶ E.4 Displaying FICON IFCC and Other errors on page 453:
  - FICON error types
  - Error detail information
  - I/O trace information
  - PPE (Purge Path Extended) information
- ▶ E.5 Displaying the Fibre Channel Analyzer panel on page 466:
  - FICON link incident log

- ▶ E.6 Displaying the System I/O Configuration Analyzer panel on page 469:
  - PCHID-Control Unit definition
  - PCHID-Partition definition
  - PCHID-Switch Link definition
  - PCHID-Node Type-Link address information

## E.1 HMC and SE information

This section details how to find HMC and SE LIC versions, user ID information, and how to logon to the HMC/SE. Ensemble user IDs are added and all of HMC/SE tasks are displayed in the Tree style user interface in this Appendix.

### E.1.1 HMC and SE versions and user IDs

The HMC and SE provide System z and FICON channel path information from a hardware point of view. All panels we discuss here are based on the following:

- ▶ The hardware being used for window captures is a z196 Hardware Management Console (HMC) and z196 Support Element (SE), both running at version 2.11.1 which corresponds to Driver level 93.
- ▶ Users are logged on in SYSPROG mode (User role) or Service mode using the Tree User Interface (UI) style depending on the tasks required.

A System z server can also be set to the Classic User Interface style. This UI style provides the exact same information, but looks different from Tree UI style. For details about this UI style or about how to change the UI style, refer to *HMC Operation Guide*, SC28-6905.

The panels can differ if you are looking at another processor generation or LIC level. However, the information and data you can gather are the same.

Table E-1 HMC and SE LIC level and machine supported is the summary of HMC and SE LIC levels and the supported machine information. We can see each LIC version information at the title panel of each HMC or SE.

- ▶ LIC level 2.11.x support z196 and z114
- ▶ LIC level 2.10.x support z10 BC and z10 EC
- ▶ LIC level 2.9.x support for z9 BC and z9 EC

*Table E-1 HMC and SE LIC level and machine supported*

| Project Code Name | Machine family | Machine type | Firmware driver | LIC version |
|-------------------|----------------|--------------|-----------------|-------------|
| zEnterprise 114   | z114           | 2818         | 93              | 2.11.1      |
| zGryphon GA1      | z196           | 2817         | 86              | 2.11.0      |
| z10 BC zMR GA2    | z10 BC         | 2098         | 79              | 2.10.2      |
| z10 EC eClipz G32 | z10 EC         | 2097         | 79              | 2.10.2      |
| z10 BC zMR GA1    | z10 BC         | 2098         | 76              | 2.10.1      |
| z10 EC eClipz GA2 | z10 EC         | 2097         | 76              | 2.10.1      |
| Pollux GA2        | z9 BC          | 2096         | 67              | 2.9.2       |
| Danu GA3          | z9 EC          | 2097         | 67              | 2.9.2       |

## E.1.2 HMC and SE default user IDs and passwords

Table E-2 HMC and SE default user IDs and passwords is the summary of user IDs and passwords necessary to perform operational and service tasks at HMC and SE. The IDs and passwords are the default at install. You have the option to modify or create additional IDs. Two more IDs (ensoperator and ensadmin) are added for Ensemble capable HMC and SE(LIC 2.11.x).

Table E-2 HMC and SE default user IDs and passwords

| Default user ID<br>*Case Sensitive | Default password<br>*Case Sensitive | Authority                        |
|------------------------------------|-------------------------------------|----------------------------------|
| sysprog                            | password                            | System Programmer authority      |
| service                            | servmode                            | Service authority                |
| acsadmin                           | password                            | Administrator authority          |
| operator                           | password                            | Basic operator authority         |
| advanced                           | password                            | Advanced operator authority      |
| ensoperator                        | password                            | Ensemble operator authority      |
| ensadmin                           | password                            | Ensemble administrator authority |

## E.1.3 Logging on to the HMC and SE

Three major components in Tree Style UI are shown at Figure E-1 on page 428:

- ▶ Navigation pane
- ▶ Work pane
- ▶ Tasks pad

To log on to the HMC and SE in Tree Style UI, do the following:

1. Log on at the HMC in sysprog or service mode depending on the tasks.
2. Click **System Managements** → **Systems** in navigation pane.
3. Highlight the CPC that you want to work with.
4. Double-click the **Single Object Operations (SOO)** task in the CPC Recovery task list in the Tasks pad.

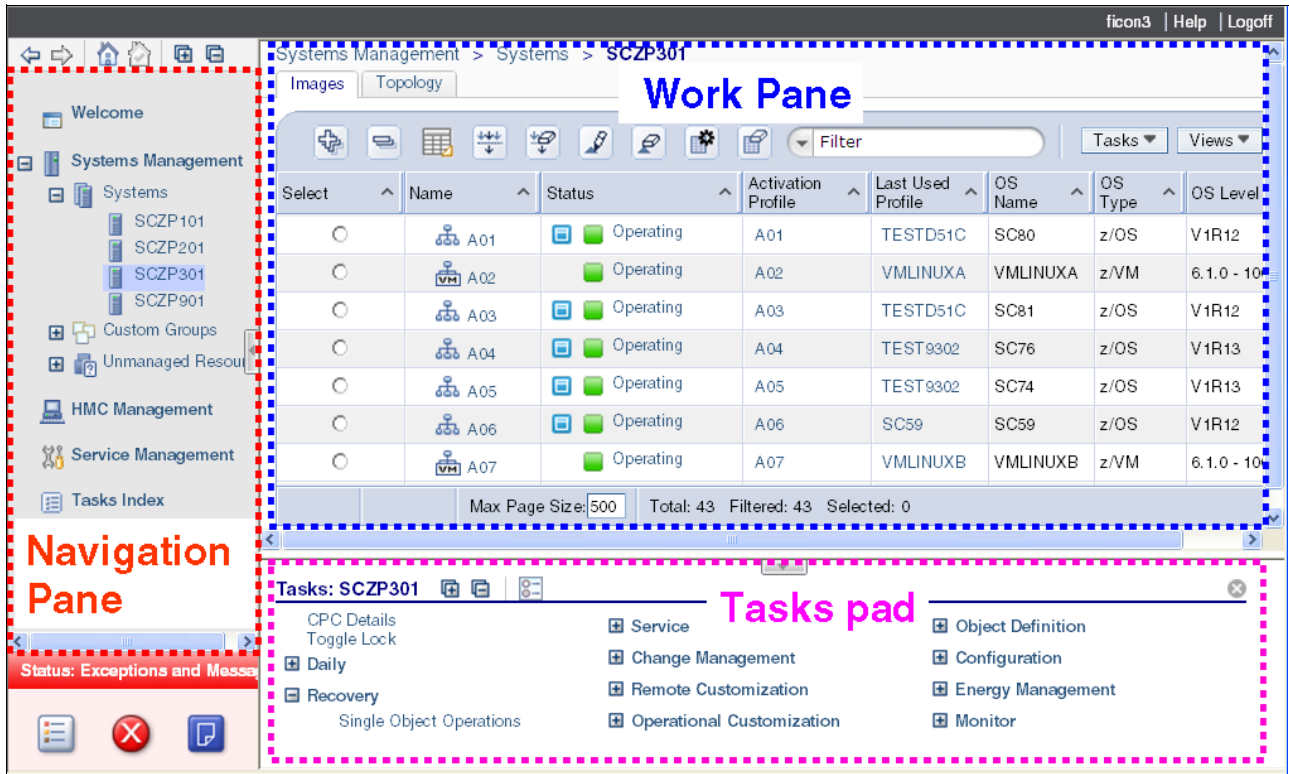


Figure E-1 HMC and SE Tree Style User Interface panel

5. Figure E-2 is SE panel after SOO performed in Tree Style User Interface.

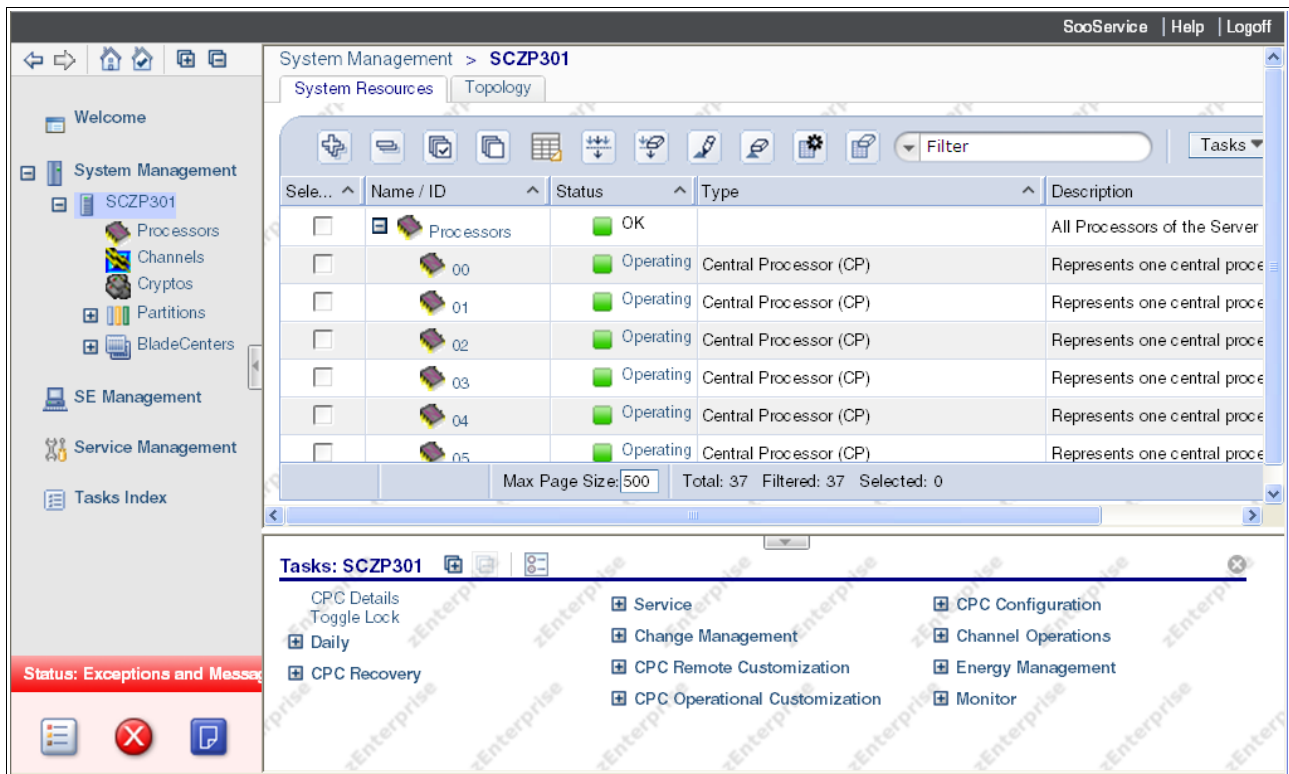


Figure E-2 SE panel after SOO

Some SE tasks are available at service mode authority only with Service Status enabled. Follow the steps to enable the Service Status at SE.

6. Highlight the CPC that you want to work with at the SE Navigation Pane.
7. Double-click **Service Status** in the Service task list at the Tasks pad.
8. Click the Object (CPC name) and click **Options**.
9. Click **Enable service status**.
10. Click **Save** and **Cancel**.

## E.2 Displaying CPC resource information

You can get CPC details and Physical resource information through the CPC Details task and View Frame Layout task HMC and SE panels.

### E.2.1 Displaying the System z CPC details panel

To open the CPC details work area:

1. Log on to the HMC and SE in sysprog or service mode (see Logging on to the HMC and SE on page 427 for more information about this topic).
2. On the SE Navigation Pane, click **System Managements**.
3. Highlight the CPC to work with.
4. Double-click **CPC Details** task in the Tasks pad.
5. Click the **Instance Information** tab in the CPC Details panel.

Figure E-3 on page 430 shows the CPC status of a CPC called SCZP301. The panel has four selectable tabs. We discuss the contents of the Instance Information and Product Information tabs.

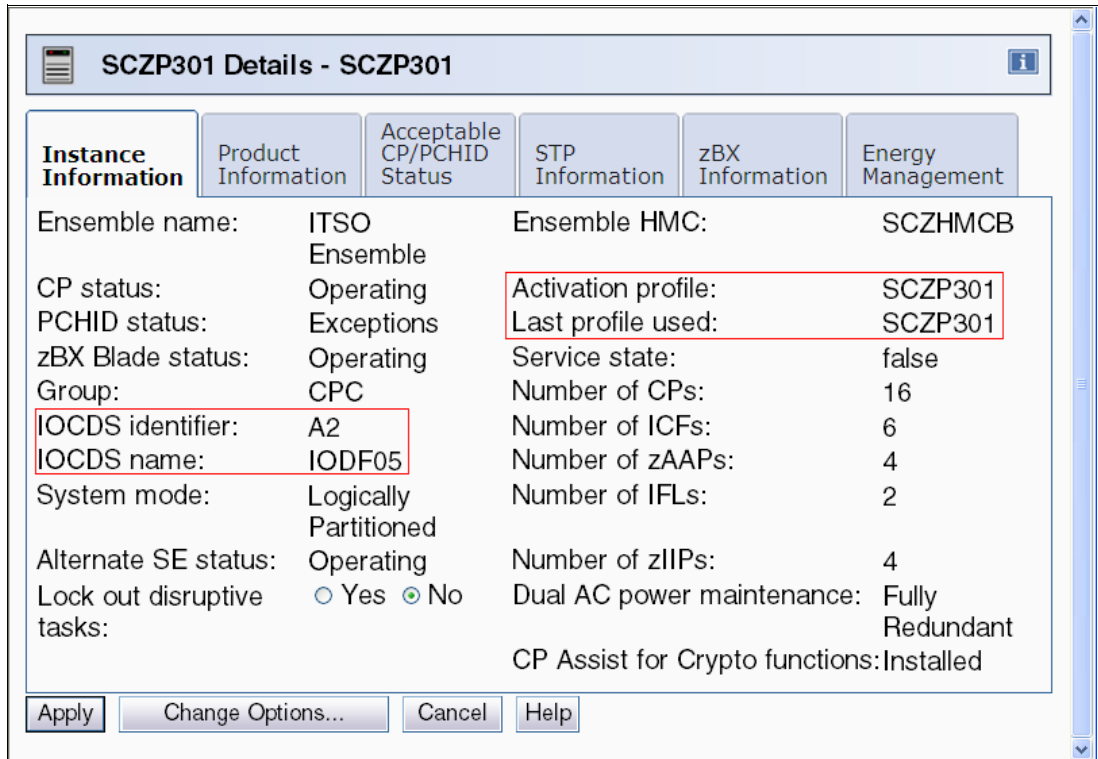


Figure E-3 CPC Details Instance Information panel

As highlighted in the CPC Instance Information panel, the following information can be retrieved (the values shown reflect our scenario):

- The IOCDS slot used to Activate the CPC: A2
- The IOCDS name used during Activation: IODF05
- The Activation Profile associated with the CEC: SCZP301
- The Activation Profile last used to Activate the CEC: SCZP301

6. Click the **Product Information** tab in the CPC Details panel.

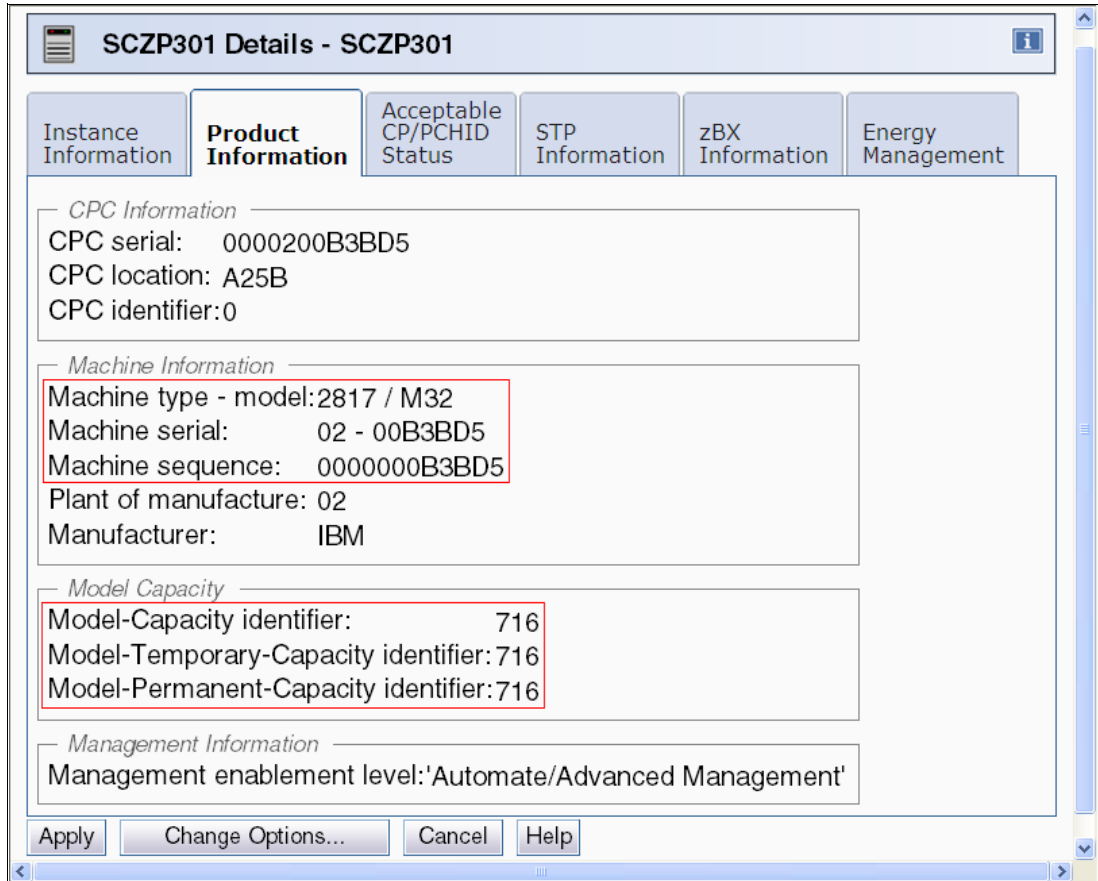


Figure E-4 CPC Details CPC Product Information

As highlighted in the CPC Product Information panel shown in Figure E-4, the following information can be retrieved from this panel:

- ▶ The machine type and hardware model: 2817 / M32
- ▶ The machine serial number: 02-00B3BD5
- ▶ The Model-Capacity Identifier: 716
- ▶ The Model-Temporary Capacity Identifier: 716
- ▶ The Model-Permanent Capacity Identifier: 716

## E.2.2 Finding a physical resource on a System z server

The following steps explain how to find a physical resource, such as a FICON channel, in the System z server. We discuss two SE panels:

- ▶ The View Frame Layout panel shows the location of the installed I/O cages where the FICON channels are normally installed.
- ▶ The Channel to PCHID Assignment panel shows important information about the PCHIDs and CHPIDs and the adaptor types and their physical locations.

### E.2.2.1 Displaying the System z server View Frame Layout panel

Follow these steps to display the View Frame Layout panel on the HMC or SE:

1. Log on to the HMC or SE in sysprog mode (see Logging on to the HMC and SE on page 427).

2. On the HMC Navigation Pane, click **System Managements**.
3. Highlight the CPC to work with.
4. Click the **Configuration task** in the Tasks pad.
5. Double-click **View Frame Layout** task in the Configuration task.

**Note:** The SE has a same panel under CPC configuration tasks that allows the Frame Layout panel to be viewed.

***Displaying the View Frame Layout panel***

Figure E-5 on page 433 shows the View Frame Layout SE panel. Clicking over the presented cages causes a “details view” to display. Selecting this option provides additional information about the frames and cages, as shown in the details at Figure E-5 on page 433. The Cage or Drawer locations are not fixed and changed depending on what type of cages or drawers and how many of them installed. We need to know each type of the installed I/O cage or drawer with its location through the View Frame Layout panel. A new I/O drawer, PCIe IO drawer is introduced to support more I/O slots in small dimension from on zEnterprise z196 and z114.

There are 3 different types of I/O cage or I/O drawer currently used for System z as shown in Table E-3 I/O Cage and I/O drawers.

*Table E-3 I/O Cage and I/O drawers*

| <b>Type</b>     | <b>Dimension</b>    | <b>I/O card slots</b> | <b>Machine supported</b> |
|-----------------|---------------------|-----------------------|--------------------------|
| I/O Cage        | 5U EIA <sup>a</sup> | 8                     | All of System z          |
| I/O Drawer      | 7U EIA              | 32                    | z196, z114, and z10      |
| PCIe I/O Drawer | 14U EIA             | 28                    | z196, and z114           |

a. U: Unit, EIA: Electronic Industry Association



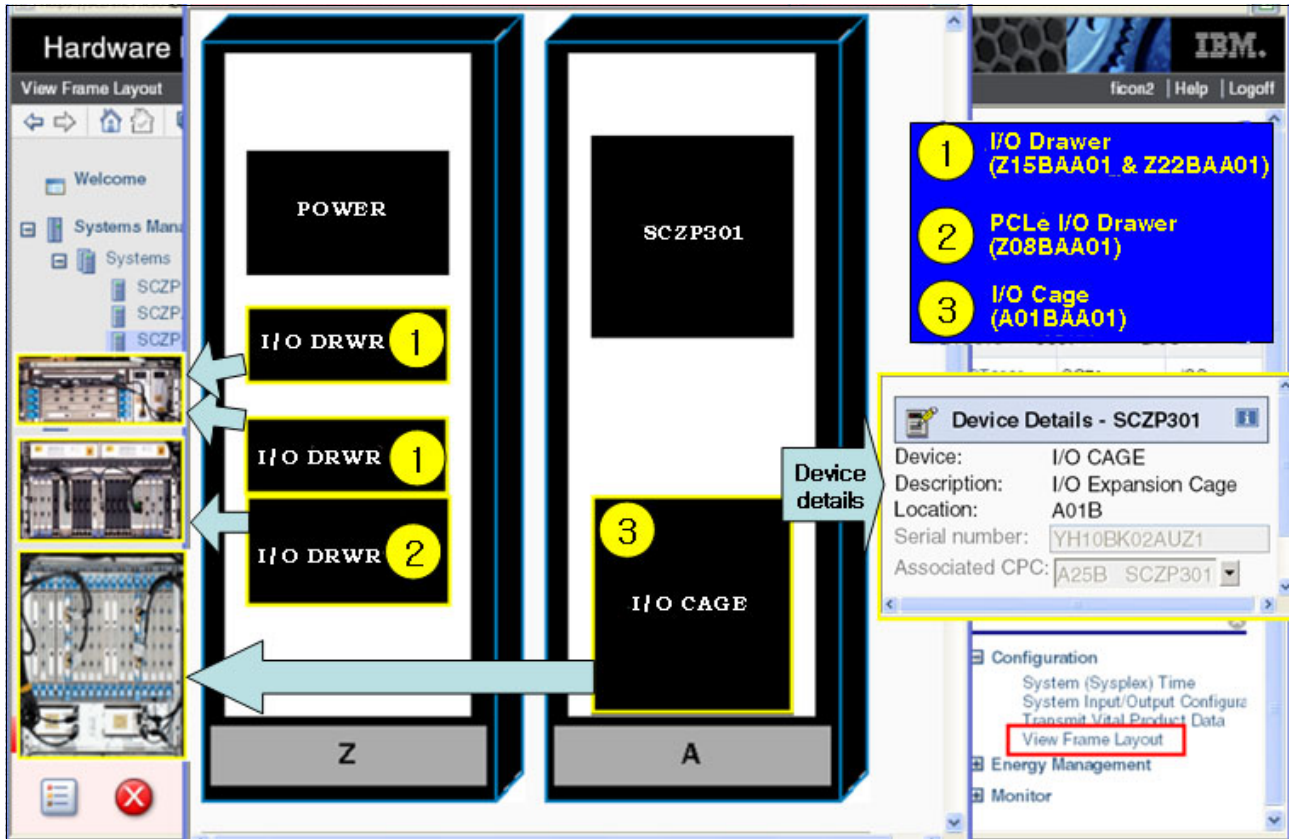


Figure E-5 HMC View Frame Layout panel

### E.2.2.2 Displaying the Channel to PCHID Assignment panel

To display the Channel to PCHID Assignment panel on the SE:

1. Log on to the HMC and SE in sysprog or service (see Logging on to the HMC and SE on page 427).
2. On the SE Navigation Pane, click **System Managements**.
3. Highlight the CPC to work with.
4. Click the **CPC Configuration** task in the Tasks pad.
5. Double-click **Channel PCHID Assignment** task in the CPC Configuration task list.

The PCHID Assignment panel appears, as shown in Figure E-6 on page 434.

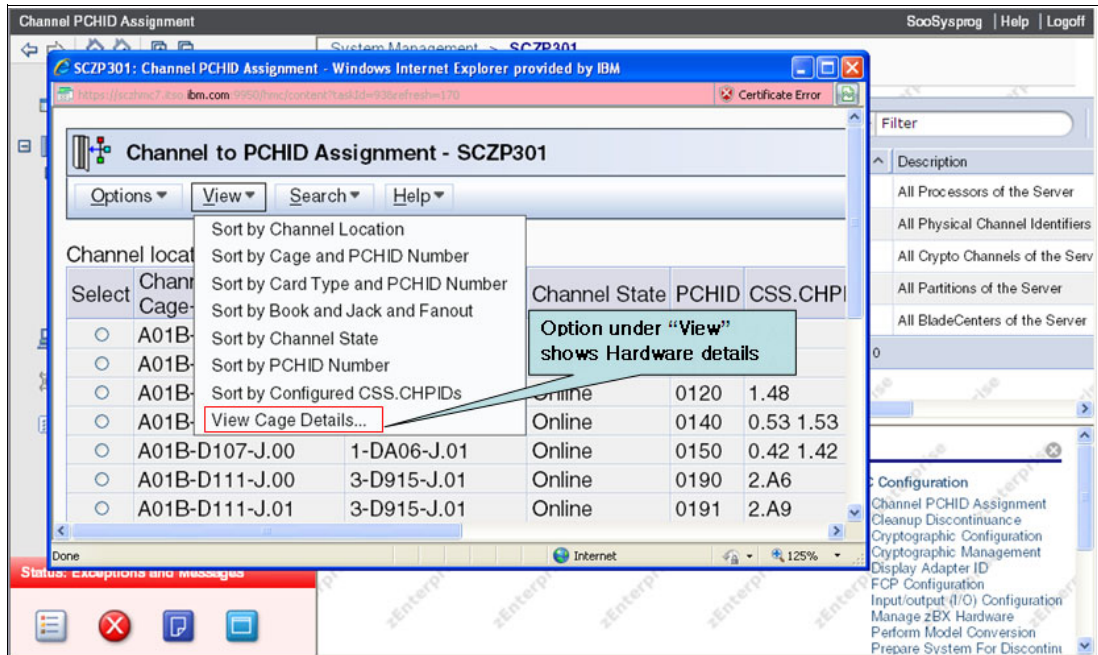


Figure E-6 Channel PCHID Assignment

In addition to showing search and different view modes, this panel displays information about the channels in use, such as:

- Channel Location
  - Physical location of a channel including Cage, Card Slot and Jack (card port)
- Book and Fanout
  - Provides information about what fanout is used to access the channel, its book location and port
- Channel State:
  - The current state of the channel, such as Online and Offline
  - PCHID: The channel's corresponding PCHID number
  - CSS.CHPID: The CHPID and CSSs that this channel belongs to
  - Card Type: The channel's corresponding hardware

6. The View Cage Details option displays PCHID and CSS.CHPID values per I/O card slot and jack for every installed I/O drawers or I/O cages, as shown at Table E-7 on page 451.

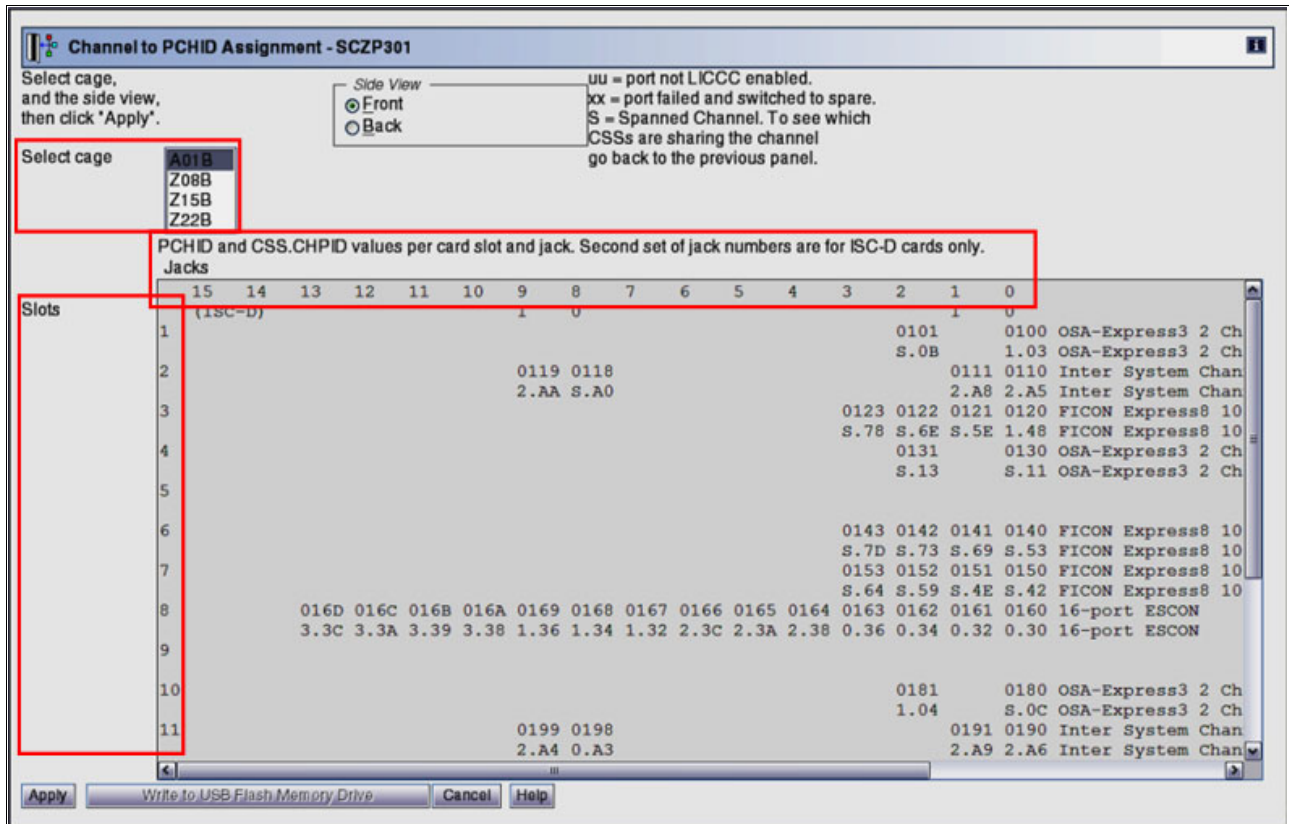


Figure E-7 View Cage Details option at Channel PCHID assignment panel


## E.3 Displaying individual FICON channel information





You can use the support element workplace to determine the state and status of specific channel paths in the input/output (I/O) configuration of the central processor complex (CPC). You can have different FICON card generations from FICON Express, Express 2, Express 4, and Express 8. The label for each channel path's icon includes its PCHID, state, and status. When you need more detailed information about determining problems, you can use the support element workplace to perform the Channel Problem Determination.

### E.3.1 FICON icon

FICON icons at the SE panel are different from each FICON card generations, as shown at the following table, Table E-4 FICON icons. From the FICON Express 2 card, the FICON icons are the same regardless of the channel definition type, FC or FCP. Refer to Table 2-1 on page 37 for the detailed FICON feature descriptions.

Table E-4 FICON icons

| ICON  | Channel type | Description     | Machines supported  |
|---|--------------|-----------------|---------------------|
|  | FC<br>FCP    | FICON Express 8 | z196, z114, and z10 |

| ICON  | Channel type | Description        | Machines supported    |
|---|--------------|--------------------|-----------------------|
|  | FC<br>FCP    | FICON<br>Express 4 | z196,z114,and z10     |
|  | FC<br>FCP    | FICON<br>Express 2 | z10 or prior machines |
|  | FC           | FICON<br>Express   | z10 or prior machines |
|  | FCP          | FICON<br>Express   | z10 or prior machines |

### E.3.2 Displaying detailed FICON channel information

To display detailed FICON channel information:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. On the SE Navigation Pane, click **System Managements**.
3. Highlight the CPC to work with.
4. Select the expand icon of the CPC, and click **Channels** at the system resources of the selected CPC.
5. All of PCHIDs are displayed at the Work Pane, as shown at Figure E-8 on page 437. Double-click the **PCHID** you want to work with.

Figure E-8 on page 437 shows the PCHID Details panel for PCHID 0120. Note that there is a panel called Status in the field. The status reflects the current state of the PCHID or CHPID.

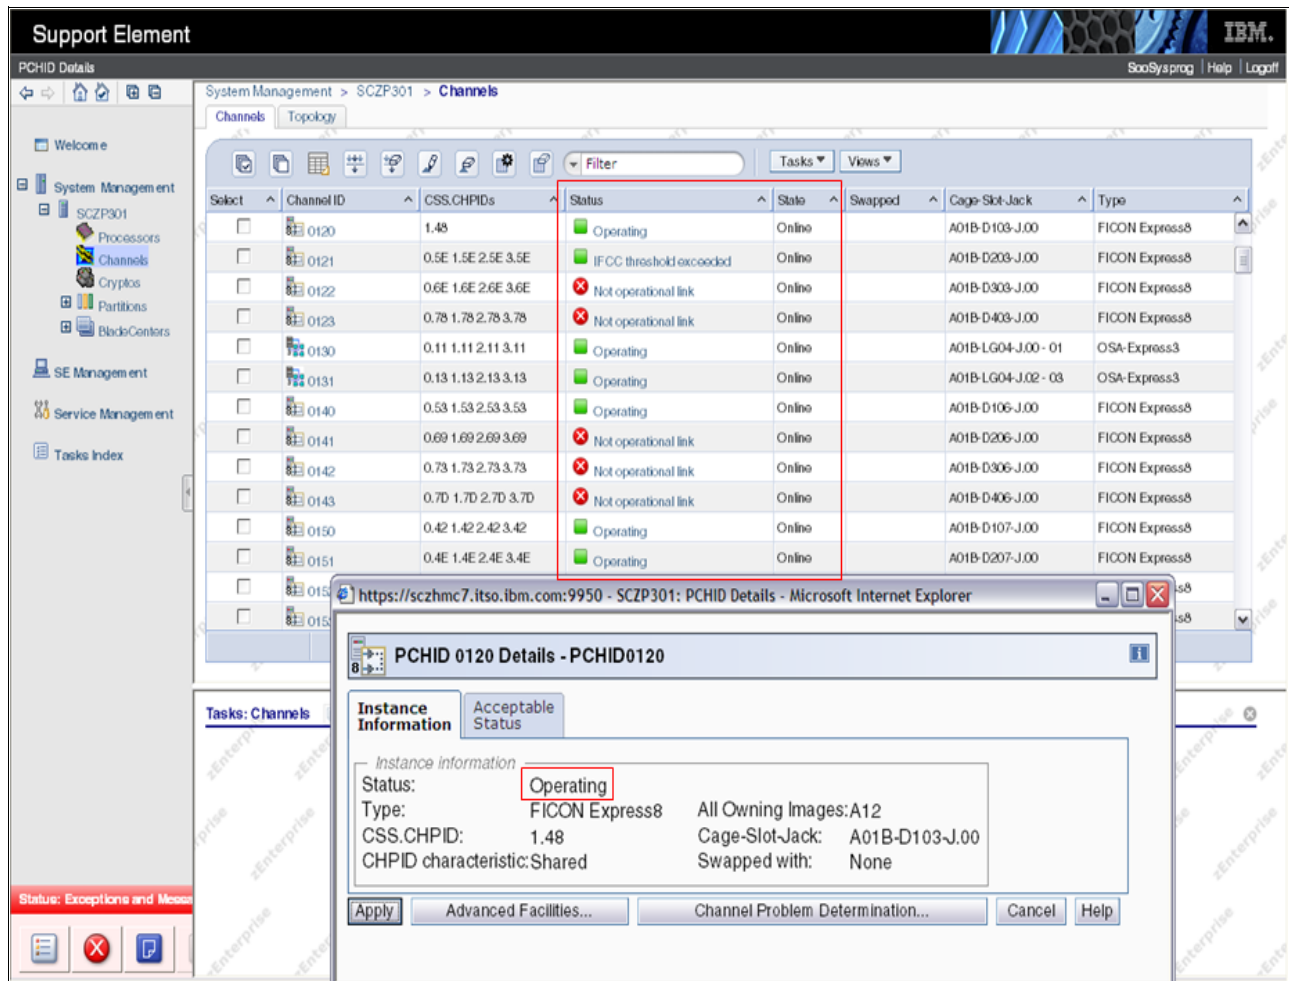


Figure E-8 PCHID 0120 Details panel

This list shows the possible state or status that a particular PCHID/CHPID might be in:

- Stopped** This state indicates that the channel is defined, but not in use.
- Operating** This is the normal indication of a working CHPID.
- Stand-by** This indicates that the CHPID has been taken offline.
- Stand-by/Reserved** This means that the channel has been put in “service” mode for repair or test.
- IFCC Threshold Exceeded** This is a channel status condition that will appear when the IFCC events have exceeded a coded preset value. For FICON channels, the threshold established limit number is 4. For each IFCC detected, this value will be decremented by the code. When a value of 0 (zero) is reached, this will be indicated in the CHPID status/state. The threshold value is shown in the Analyze Channel Status SE PD panel.
- Definition Error** This condition is caused by the defined device not matching the attached device. The channel is a FICON channel but it is not compatible with the channel type defined in the IOCDs.

**Note:** When a definition error is detected, the CHPID icon will default to the icon that represents a FICON Converter CHPID, and it is not necessarily representative of the defined channel type.

**Sequence Time Out** A FICON “sequence” is a FICON frame or a group of FICON frames transmitted between the Channel and the CU or between the CU and the channel. The number of frames that make up a sequence is determined by specific fields inside the frames themselves. Sequences are logically grouped to form an exchange. An exchange represents an I/O operation in the FICON implementation. A Sequence Time Out is detected by the channel or the CU when an ongoing exchange does not detect the arrival of a required sequence.

**Sequence Not Permitted** This error is reported in the Analyze Channel information panel as “Illegal Sequence.” It reports that the FICON channel-attached control unit or FICON Director port does not support the presented FICON service parameters.

Table E-5 FICON CHPID and switch port link error logging chart is the summary of FICON CHPID status related to the attached Switch port status and Switch incident log for FICON channel PD based on *z196 Maintenance Guide, SY28-6892-02*.

Table E-5 FICON CHPID and switch port link error logging chart

| FICON CHPID status  | Link error  | Switch port status      | Switch incident log                             |
|---|---|-------------------------|---|
| Online/Sequence time out                                  | Outbound loss of light from FICON CHPID             | Link Failure (LinkFail) | Loss of Signal (LOSS)                           |
| Online/Loss of Signal                                     | Inbound loss of light to FICON CHPID                | Link Failure (LinkFail) | Loss of Signal (LOSS)                           |
| Online/Operating or IFCC threshold                        | Outbound bit error of FICON CHPID                   | Normal Operating status | Nothing   |
| Online/Operating or Bit error threshold or IFCC threshold | Inbound bit error to FICON CHPID                    | Normal Operating status | Nothing   |
| Online/Operating  | Outbound loss of light from switch destination port | Link Failure (LinkFail) | NOS(Not Operation Sequence) or Sequence timeout |
| Online/Operating  | Inbound loss of light to switch destination port    | Link Failure (LinkFail) | LOSS(Loss of Signal)                            |
| Online/Operating  | Outbound bit from switch destination port           | Normal Operating status | Nothing   |
| Online/Operating  | Inbound bit error to switch destination port        | Normal Operating status | Bit error threshold                             |

### E.3.3 Displaying the channel problem determination task panels

When you need more detailed information about determining problems, you can use the support element workplace to perform *channel problem determination*. Perform channel

problem determination to get the following types of information, referred to as problem determination information, for a channel path:

- ▶ Analyze channel information
- ▶ Analyze subchannel data
- ▶ Analyze control unit header
- ▶ Analyze paths to a device
- ▶ Analyze device status
- ▶ Analyze serial link status
- ▶ Display message buffer status(Not applicable for FICON channel)
- ▶ Fabric login status(Not applicable for FICON channel)
- ▶ Analyze link error statistics block
- ▶ Optical Power Measurement (Service mode only)

Some information is available in service mode with the Service status enabled. For instance, when you need to display OPM (Optical Power Measurement), it is available in service mode only, as shown at Figure E-9.

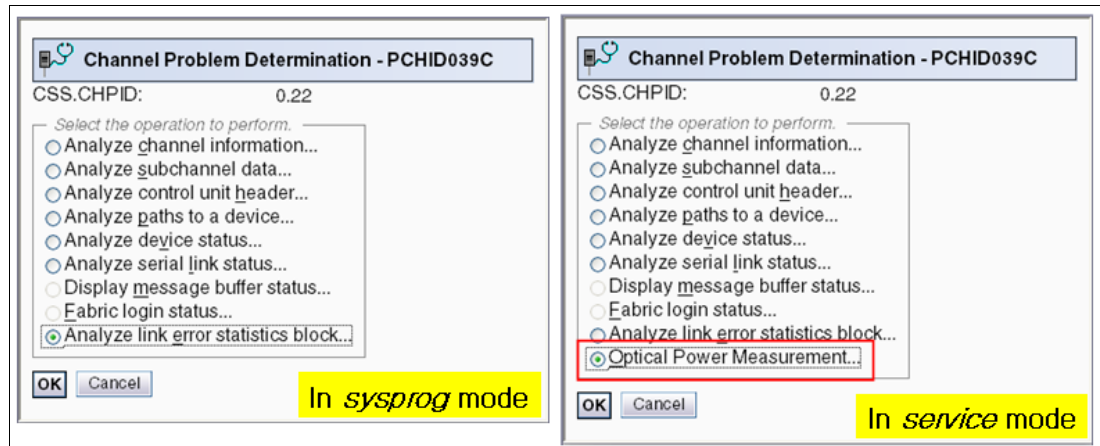


Figure E-9 Channel PD panels in sysprog mode and service mode

### E.3.3.1 Displaying Channel Problem Determination task panel on SE

To display the Channel Problem Determination task, log on to the HMC and SE (see Logging on to the HMC and SE on page 427) in sysprog or service mode depending on the channel information required.

There are four different ways to display the Channel Problem Determination task. You can display it through one of four steps (1, 2, 3, or 4 steps):

1. On the SE Navigation Pane, click **System Managements**.
2. Highlight the CPC to work with. On the SE Work pane, click **Channels** under the Name/ID column:
  - a. Click a PCHID icon to work with.
  - b. Click **Channel Problem Determination** at the CPC details panel.
  - c. Select a Partition and CSS.PCHID combination, and press **OK**.
3. On the SE Navigation Pane, click the **CPC expand** icon, and select **Channels** under the CPC expand icon:
  - On SE Work Pane, click a **PCHID** icon to work with under Channel ID column.
  - Click **Channel Problem Determination** at the CPC details panel.
  - Select a Partition and CSS.PCHID combination, and click **OK**.

4. On the SE Navigation Pane, click the **CPC expand** icon, and select **Channels** under the CPC expand icon:
  - a. On SE Work Pane, Select a PCHID to work with under Select column.
  - b. Click a Channel Problem Determination task at the CHPID operation task in Tasks pad.
  - c. Select a Partition and CSS.PCHID combination, and click **OK**.
5. On the SE Navigation Pane, click the **CPC expand** icon, and select **Channels** under the CPC expand icon:
  - a. On the SE Work Pane, select a PCHID to work with under the Select column.
  - b. Click a Channel Problem Determination task at the Channel operation task in Tasks pad.
  - c. Select a Partition and CSS.PCHID combination, and click **OK**.

**Note:** You can display the same Channel Problem Determination panel with CHPID instead of PCHID, when you select **Partition** object under the CPC resource expand at the SE Navigation pane.

### E.3.3.2 Displaying the Analyze Channel Information panels

To open the Analyze Channel Information (PD) panels:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. Display the Channel Problem Determination panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Channel Information**, and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

The Analyze Channel Information panel is shown in Figure E-10.

| Analyze Channel Information - CHPID1.21 |  |
|---|--|
| Channel type: FICON Native              | Hardware type: E0                      |
|   | Hardware subtype: 04                   |
| Partition ID: 11                        | 2 byte control unit                    |
| MIF image ID: 1                         | link address defined: Yes              |
| Channel mode: Spanned                   | Absolute address: 00000009F4E8800      |
| CHPARAM: 00                             |  |
| CSS.CHPID: 1.21                         |  |
| PCHID: 0399                             |  |
| Switch number: 65                       |  |
| Switch number valid: 1                  |  |
| State: Online                           | IFCC threshold: 4                      |
| Status: Operating                       | Channel link address: 659200           |
| Image chnl state: Online                | Temp error threshold: 4                |
| Image chnl status: Operating            | Suppress: 0000000000000000             |
| Error code: 00                          | SAP Affinity: 05                       |
| Ber inbound: 0                          | Card description: FICON Express8       |
| Ber outbound: 0                         | 10km LX                                |
|   | Connection rate: 8Gbps                 |
| Node type: Self                         | Node type: Attached                    |
| Node status: Valid                      | Node status: Valid                     |
| Flag/parm: 10000721                     | Flag/parm: 00200A92                    |
| Type/model: 002817-M32                  | Type/model: SLKWRM-DCX                 |
| MFG: IBM                                | MFG: BRD                               |
| Plant: 02                               | Plant: CA                              |
| Seq. number: 0000000B3BD5               | Seq. number: 1AFX0642C00Z              |
| Tag: F021                               | Tag: 6540                              |
| World wide node name: 5005076400CB3BD5  | World wide node name: 100000051E439201 |
| World wide port name: 5005076028BF0F58  | World wide port name: 204000051E439201 |

Figure E-10 Analyze Channel Information PD panel for CHPID 1.21



Note the blocks that are highlighted:

- ▶ On the top left side, in the block labeled *Definition Information*, notice that *Partition ID 11* is the LPAR ID associated to the partition where this panel was selected from. *MIF image ID 1* corresponds to the LPAR MIF ID on *CSS 1*; also note that *CHPID 21* is reported as *Spanned*, meaning that it has been shared between LPARs in different CSSs and is associated to *PCHID 0399*. It also reports that *CHPID 21* is connected to *Switch number 65*.
- ▶ The block labeled *Status Information* indicates the status and the state information of this CHPID. This block also displays the *Error code* and the *Bit Error Counter (Ber)* accumulated values for the CHPID. The *Error code*, when not equal to 0 (zero), might be displayed using a panel button called *Error details*.
- ▶ The block on the top right side labeled *Hardware Card type Cascade definition* indicates that this channel was defined using *2 byte control unit link address*. Normally, this is an indication that cascaded switch topology is being used.

**Note:** HCD allows a dual-byte address definition even when a single FICON Director is used in the configuration. (In this case, the entry port and the destination port belong to the same switch number).

- ▶ The block labeled *Error Status* shows the link status, the connection port, and the *SAP affinity* associated to the CHPID. The *IFCC threshold* has an initial value of 4 and decrements at every occurrence of an IFCC on this channel. When the value reaches 0 (zero), the indication of a *IFCC Threshold Exceeded* will be associated to the CHPID icon on the SE desktop. The *Temporary error threshold* works the same way as the *IFCC threshold*. The *SAP affinity* indicates which of the SAPs this CHPID is associated to.
- ▶ The block labeled *Card Type Connection rate* shows the *FICON card description type* associated with this CHPID. It has a direct relationship with the information provided about the *Hardware type* and *Subtype* presented in the top right corner of the panel.

### Understanding Node (ND) information displayed on this panel

The data listed here is received from the request node identification (RNID) during the initialization of the link:

|                        |  |
|------------------------|--|
| <b>Node type</b>       | <i>Self</i> means the shown Node Descriptor (ND) information is from the channel itself.<br><i>Attached</i> indicates the shown Node Descriptor (ND) information is from the attached node.  |
| <b>Node status</b>     | <i>Valid</i> means there has been a successful port login.<br><i>Invalid</i> means there has not been any port login.<br><i>VALID BUT NOT CURRENT</i> means there has been a port login, but the link got dropped afterwards.                        |
| <b>Type/Model</b>      | This field indicates the machine type and model of the node.   |
| <b>Plant</b>           | This field indicates the plant where the node was manufactured.  |
| <b>Sequence number</b> | This field represents the serial number of the node.   |
| <b>Tag</b>             | The <i>Tag</i> field represents the <i>CSS.CHPID</i> for System z server. For the attached node, you need to know whether the tag is vendor-dependent, and you must also refer to the vendor's documentation to understand how to decode this field. |

The System z server adds a new flavor to the Tag field, as explained here:

The first character of the Tag field represents the CSS. The last two digits of the Tag field represent the CHPID (the second digit showing a zero (0) is reserved).

Table E-6 CSS-to-tag relationship illustrates the CSS/Tag relationship. The tag value in the top line gives the corresponding CSS IDs that the node has been defined to. For example, with Tag = F021, the F indicates that the node is defined to CSS0, CSS1, CSS2, and CSS3 represents CHPID 21. Table E-6 CSS-to-tag relationship lists the CSS-to-tag relationship.

Table E-6 CSS-to-tag relationship

| CSS\Tag | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | A | B | C | D | E | F |
|---------|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0       |   |   |   |   |   |   |   | X | X | X | X | X | X | X | X |
| 1       |   |   |   | X | X | X | X |   |   |   |   | X | X | X | X |
| 2       |   | X | X |   |   | X | X |   |   | X | X |   |   | X | X |
| 3       | X |   | X |   | X |   | X |   | X |   | X |   | X |   | X |

**World Wide Node Name** This is a unique 64-bit address that is used to identify the node in an FC topology. This address is assigned by the manufacturer.

**World Wide Port Name** This is a unique 64-bit address that is used to identify the port in an FC topology. This address is assigned by the manufacturer.

### E.3.3.3 Displaying the Analyze Subchannel Information PD panel

To open the Analyze Subchannel Information panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. Display the Channel Problem Determination panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Subchannel Information**, and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

4. Select the option **Search by Device Number** (you can also search by a unit address or subchannel number, if appropriate), and click **OK**.
5. Enter the device number (or unit address or subchannel number, depending on what you selected in the previous step), and select **OK**. In our example, we used device number 9000 (see Figure E-11 on page 443).




Figure E-11 Analyze Subchannel Information PD panel

Note the fields that are highlighted:

- ▶ The field Irpt parm: 00EA0268 displays this subchannel's associated UCB SCP absolute address. During IPL device mapping process, the subchannel is connected to its UCW and this field contains the respective pointer to the UCB in SCP memory.
- ▶ The Pathing information is displayed in the block in the center of the panel. The list of candidate CHPIDs is shown in the top right corner.

This Pathing information has various fields that are bit-position-relative to the defined channel paths to this device. Each bit position represents one channel path.

Bit 0, or the first position bit from the left, corresponds to CSS.CHPID0.PCHID.

Bit 1, or the second position bit from the left, corresponds to CSS.CHPID1.PCHID, and so on.

The following fields are shown:

CHPID - Associated CHPID number.

LPM - Logical Path Mask: The bits show which CHPIDS can be used to access the device.

LPE - Logical Path Established: Indicates the paths to this device that successfully went through the pathing process.

PNOM - Path Non-Operational Mask: This field indicates which paths from the list of candidate CHPIDs are in the Not Operational state.

LPUM - Last Path Used Mask: This field indicates the last path used to access the device when successful or any abnormal condition occurred for the I/O operation. It is updated when the Ending Status is received from the device.

PIM - Path Installed Mask: This field indicates the channel paths that are defined in the IOCDS.

POM - Path Operational Mask: This field contains the paths online to the device. Initially (at System Reset), this value is set to FF (all ones) and the corresponding bit will be turned OFF when the first attempt to access the device fails (Inoperative).

PAM - Path Available Mask: This field indicates the Physically Available channel paths. The initial value will be the paths defined in the IOCDS for the LPAR.

This bit will be turned off when:

- The corresponding CHPID is configured OFF from the HMC.
- The **CF CHP(xx)**, **OFFLINE** z/OS command is issued from the operating system console.

#### **E.3.3.4 Displaying the Analyze Control Unit Header PD panel**

To display the Control Unit Header Problem Determination panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. Display the Channel Problem Determination panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Control Unit Header**, and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

4. Select the option **Search by Device Number** (you can also search by a unit address or subchannel number if appropriate), and click **OK**.
5. Enter the device number (or unit address or subchannel number, depending on what you have selected in the previous step), and select **OK**.

The Control Unit Header PD panel is shown in Figure E-12 on page 445.

| Analyze Control Unit Header - CHPID1.21  |                  |               |          |                   |             |             |              |     |       |              |      |
|--|------------------|---------------|----------|-------------------|-------------|-------------|--------------|-----|-------|--------------|------|
| CSS ID:  | 1                | CU-hdr no:    | 002E     | Sw busy mask:     | 00000000    |             |              |     |       |              |      |
| Subchnl set:   | 0                | CU type:      | 10       | Ded alleg mask:   | 00000000    |             |              |     |       |              |      |
| Absolute address:  | 0000000140103000 | CU busy mask: | 00000000 |                   |             |             |              |     |       |              |      |
| Subch no:  | 1560             | Unit address: | 00       | Shared busy mask: | 00000000    |             |              |     |       |              |      |
| Device No:   | 9000             | MIF image ID: | 1        | Init incomp:      | 00000000    |             |              |     |       |              |      |
| Note: If data is not available or the path is not valid, then those table entries will be blank. |                  |               |          |                   |             |             |              |     |       |              |      |
| Path   | CHPID            | Link          | CUADD    | Node Type         | Node Status | Flag / Parm | Type / Model | MFG | Plant | Seq. Number  | Tag  |
| 0  | 40               | ..1B..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0000 |
| 1  | 41               | ..1B..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0100 |
| 2  | 42               | ..10..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0200 |
| 3  | 43               | ..10..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0300 |
| 4  | 21               | 6641..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0133 |
| 5  | 25               | 6611..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0333 |
| 6  | 22               | 6534..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0033 |
| 7  | 26               | 6562..        | 00       | Device            | Valid       | 00310100    | 002107-932   | IBM | 75    | 0000000L3001 | 0233 |

Figure E-12 Analyze Control Unit Header PD panel

The starting point to obtain this data is provided by the device number entered when selecting this PD action. The output presented shows:

- Path** This indicates all possible paths to get to the entered device.
- CHPIDS** This indicates the channels associated to the displayed paths.
- Linkqq** This indicates the destination link address in use by each of the CHPIDs and respective paths.
- Node Type** This indicates the type of attachment to the CHPIDs and Paths.
- Node Status** This indicates the current status of the node.

**Type/Model, MFG, Plant, Seq. Number and Tag**

All together, these represent the Control Unit RNID (Remote Node Identification) information.

All the possible paths to get to the entered device are displayed, along with the CHPIDs associated to these paths and their respective FICON Director links, where applicable. This device's control unit characteristics are also displayed, along with the tags (which represent the attaching Host Bay Adapters) that the channels connect to.

In our example, we highlighted the FICON CHPID 1.21 in the control unit header frame, which is connected to a 2107 using tag 0133. Refer to DS8800 port layout for the logical port layout of a 2107, see Appendix D, Configuring the DS8000 for FICON on page 415.

Using Table D-1 on page 419 to translate tag value 0133 to a physical port results in the following:

- ▶ Rack 1 (R1)
- ▶ I/O enclosure 2 (I2)
- ▶ Card 4 (C4)
- ▶ Port 3 (T3)

**E.3.3.5 Displaying the Analyze Paths to a Device PD panel**

To display the Analyze Paths to a Device panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).

2. Display the Channel Problem Determination panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Paths to a Device** and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

4. Select the option **Search by Device Number** (you can also search by a unit address or subchannel number if appropriate), and click **OK**.
5. Enter the device number (or unit address or subchannel number, depending on what you have selected in the previous step), and click **OK**.

The Analyze Path to a Device panel is shown in Figure E-13.

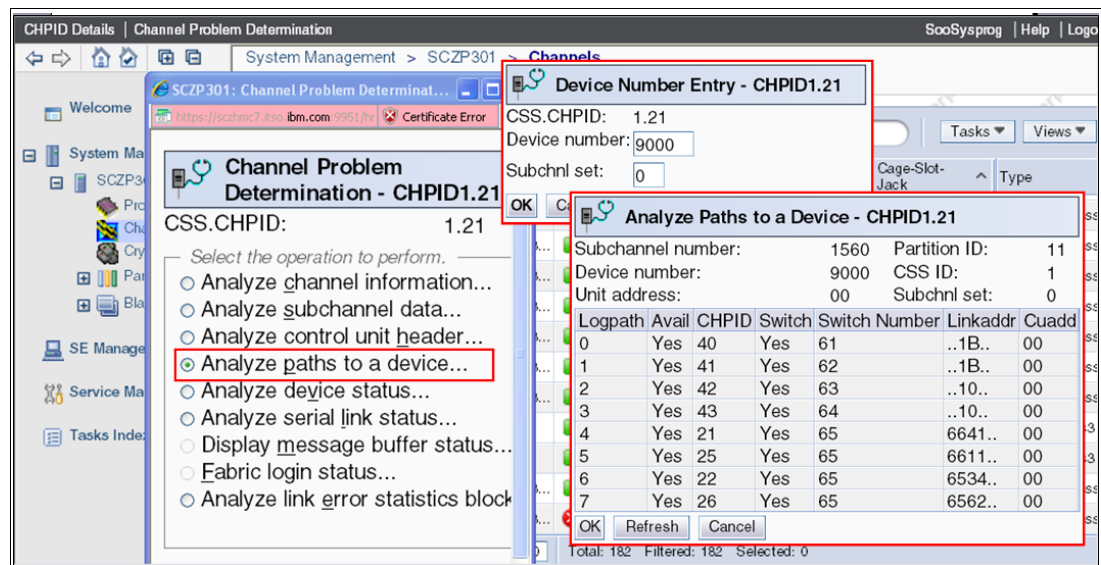


Figure E-13 Analyze Paths to a Device PD panel

It can display important device path information, as follows:

- ▶ The device's associated subchannel.
- ▶ All paths by which the device can be accessed (this is relevant only for the partition the channel path has been selected from), with the following additional information:

|                      |   |
|----------------------|---|
| <b>Avail</b>         | This shows if that path is available at that instance.                          |
| <b>CHPID</b>         | This is the associated CHPID number.  |
| <b>Switch</b>        | This shows if there is a FICON Director in the link.                            |
| <b>Switch Number</b> | This is the FICON Director number (if applicable).                              |
| <b>Linkaddr</b>      | This is the destination link address (the port address the CU is connected to). |
| <b>Cuaddr</b>        | This is the logical CU that the device is connected to.                         |

### E.3.3.6 Displaying the Analyze Device Status PD panel

To open the Analyze Paths to a Device panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).

2. Display the Channel Problem Determination panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Device Status** and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

4. Select a **Link and Control Unit Address** (Select a Switch and its link address for a Device attached to work with), and click **OK**. The Analyze Device Status PD panel is shown in Figure E-14.

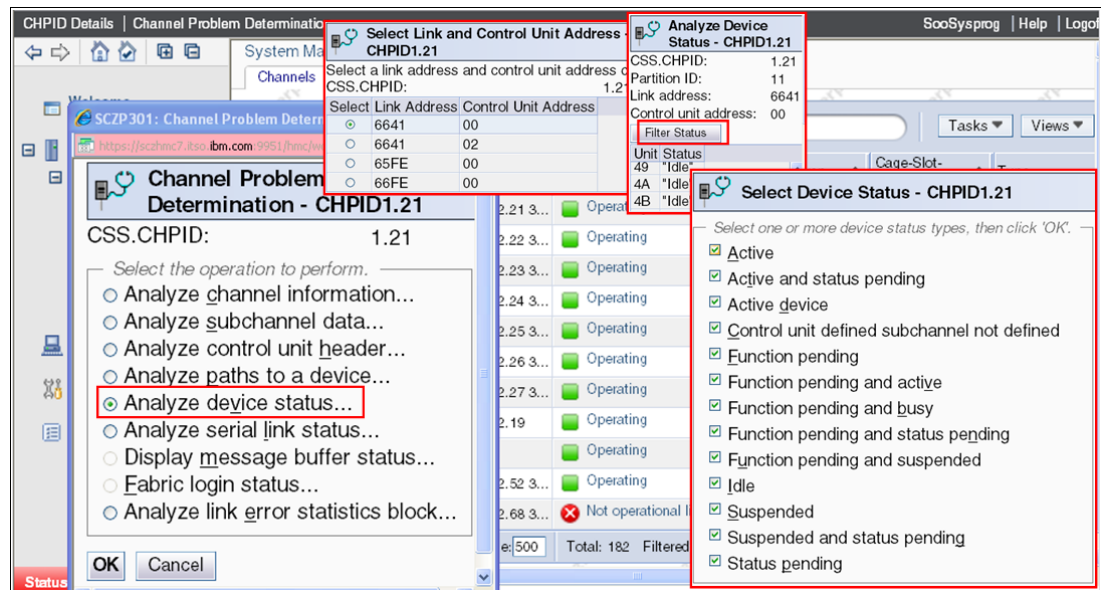


Figure E-14 Analyze Device Status panel

The Analyze Device Status panel reports the current state of all devices pertaining to a Logical Control Unit. Note that a Filter Status button is provided to find a specific status a particular device.

### E.3.3.7 Displaying the Analyze Serial Link Status PD panel

To open the Analyze Paths to a Device panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. Display the **Channel Problem Determination** panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Serial Link Status**, and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

The Analyze Serial Link Status panel is shown in Figure E-15 on page 448.

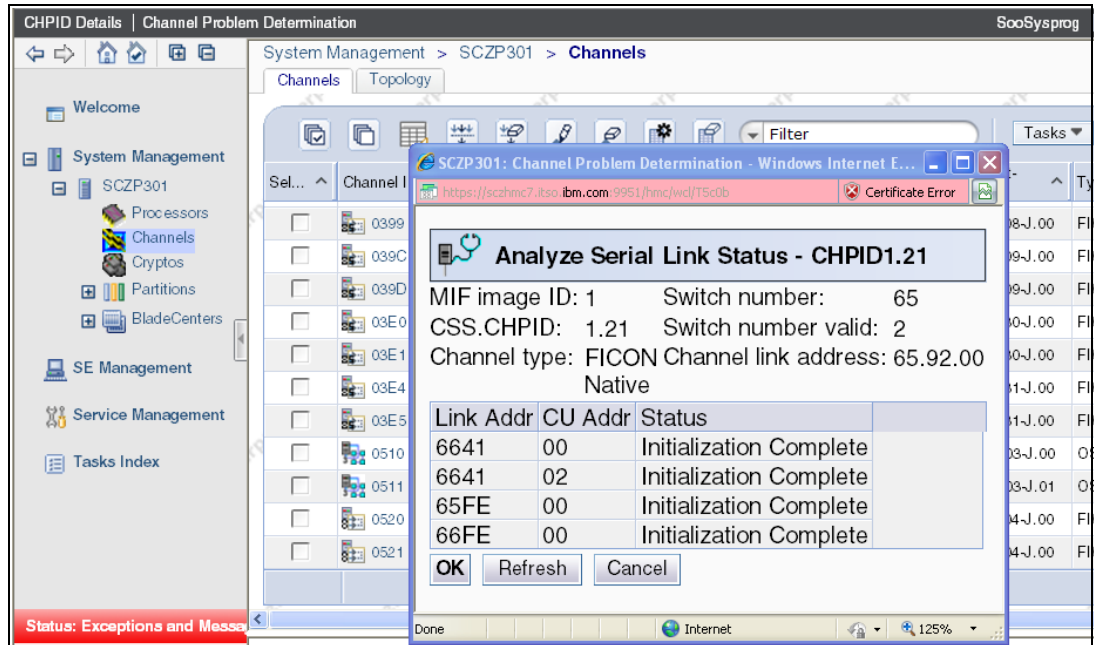


Figure E-15 Analyze Serial Link Status PD panel

The following logical pathing information is displayed:

- ▶ MIF Image ID: This is the MIF ID of the image we selected the channel from.
- ▶ CHPID: This is the logical channel consisting of CSS and CHPID.
- ▶ Channel type: This is the channel type as defined in the IOCP.
- ▶ Switch number: This is the switch number of the FICON Director that the channel is connected to.
- ▶ Switch number valid: This is the total number of switches in the links.
- ▶ Channel link address: This is the link address (port address) that the channel is connected to.
- ▶ The status of all logical link addresses can be seen in this panel, as well. This provides information about the logical CU initialization from a channel point of view. The given CU address cannot be accessed if the status is other than INITIALIZATION COMPLETE.
- ▶ Channel 1.21 from MIF Image 1 connected to Switch number 65 port address 92.
- ▶ A series of logical Control Units (residing in the same physical CU) being accessed from Channel 1.21 show the status INITIALIZATION COMPLETE.
- ▶ Any status other than INITIALIZATION COMPLETE indicates a potential problem that might need to be investigated.

Note, however, that a status of INITIALIZATION PENDING is not necessarily an error indication. It depends on whether the logical CU is expected to be available.

For example, if a channel is trying to establish a FCTC connection to another server or image but the LPAR is not activated, this would show a status of INITIALIZATION PENDING. However, the status is expected to change after the target LPAR is activated.

Other typical error messages are:

- Initialization Complete - CU Busy
- Initialization Complete - Dynamic Switch Port Busy
- Initialization in Progress



- CU Reset in Progress
- Remove Logical Path in Progress
- Link Configuration Problem - Port Reject - Address not Valid
- Switch Port Not Defined - Port Reject - Undefined Destination Address
- Switch Port Malfunction - Port Reject - Port Destination Malfunction
- Switch Port Not Available - Port Reject - Port Intervention Required
- Link Configuration Problem - Link Reject - Unrecognized Device Level
- Link Configuration Problem - Link Reject - Uninstalled Link Control Function
- Link Level Error - Link Reject - Transmission Error
- Link Level Error - Link Reject - Destination Address not Valid
- Link Level Error - Link Reject - Acquire Address Error
- Link Configuration Problem - Initialization Failure - Channel/CU mismatch
- CU Resources Exceeded - Init Failure - No Resources Available
- Link Level Error - Channel Detected Error
- Internal CTC Initialization Complete
- Link Level Error - FICON FLOGI ELS Error
- Link Level Error - FICON PLOGI ELS Error
- Link Level Error - FICON RNID ELS Error
- Link Level Error - FICON SCR ELS Error
- Link Level Error - FICON LIR ELS Error
- Link Level Error - FICON Invalid Attachment
- CU resources exceeded: This indicates that the CU ran out of logical CUs.
- CU reset in progress: This indicates that there is a reset initiated by the CU itself in progress.
- Channel/CU mismatch: This indicates that the logical CU address (defined in the CU) does not match its definition in the IOC.P
- PLOGI error: This indicates that the link initialization did not succeed to the Destination Port.

**Note:** The latest implementation of the System z Server SE Problem Determination facility improves the PLOGI and FLOGI error messages substantially. With this new implementation the panel will show, in plain English, what went wrong during the CU link initialization. Figure E-16 on page 449 shows one of the new messages issued when initialization failures occur.

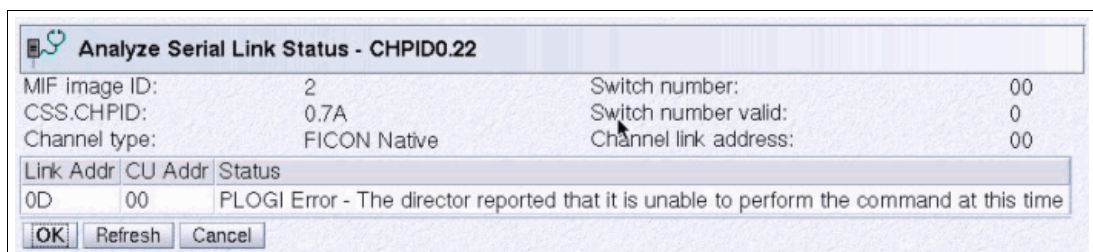


Figure E-16 New PLOGI Initialization message

### E.3.3.8 Displaying the Analyze Link Error Statistics Block PD panel

To open the Analyze Paths to a Device panel:

1. Log on to the HMC and SE (see Logging on to the HMC and SE on page 427).
2. Display the **Channel Problem Determination** panel (see Displaying Channel Problem Determination task panel on SE on page 439).
3. On the Channel Problem Determination, select **Analyze Link Error Statistics Block** and click **OK**.

**Note:** The SE will post a panel for an Image query if you select a PCHID instead of a CHPID.

The Analyze Link Error Statistics Block PD panel is shown in Figure E-17.

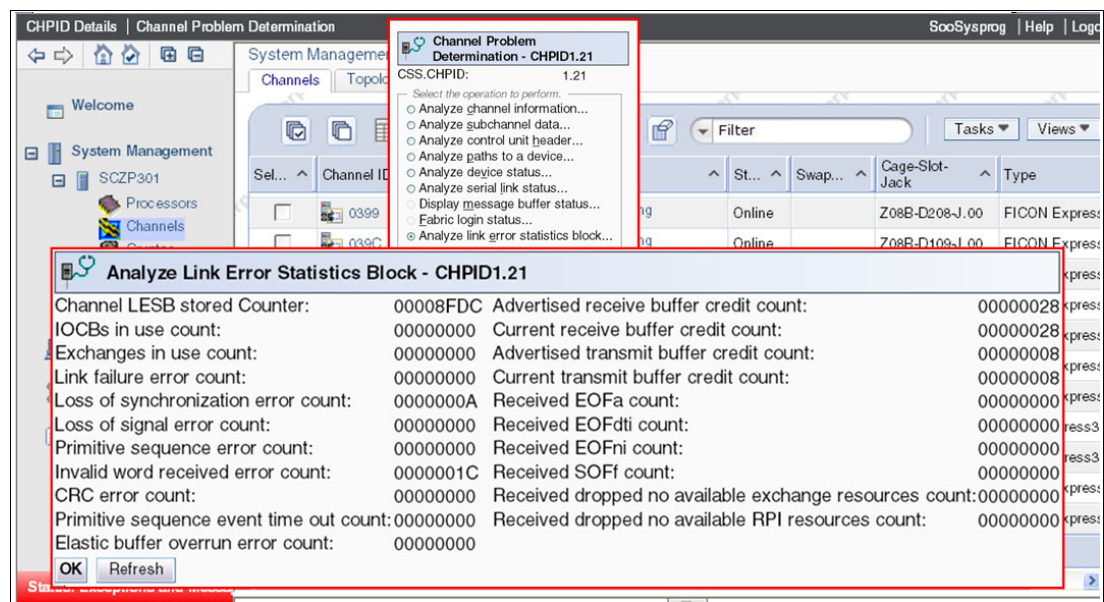


Figure E-17 Link Error Statistics Block PD panel

The Analyze Link Error Statistics Block (LESB) is a new implementation on the System z Server's Problem Determination SE panels. It displays the Channel N-Port LESB concurrently with channel operations.

The various types of errors are accumulated to provide you with a view of how the channel is operating. A refresh provides a dynamic display of the counters.

The information gathered from the Channel N-Port is also recorded in the EREP SLH records when a Purge Path Extended (PPE) operation is performed. For more information about the PPE process, refer to 10.6 FICON Purge Path Extended on page 357 and E.4.7 Purge Path Extended Log on page 461.

### E.3.3.9 Displaying the Optical Power Measurement panel

To open the Optical Power Measurement panel:

1. Log on to the HMC and SE in service mode and enable service status (see Logging on to the HMC and SE on page 427).
2. Display the **Channel Problem Determination panel** (see E.3.2 Displaying detailed FICON channel information on page 436).

3. Select the option **Optical Power Measurement**, and click **OK** at the Channel Problem Determination panel.

**Two Versions of Optical Power Measurement panels:** Depending on the System z family and the driver level you are working on, there are two versions of the panel:

- ▶ z10 with Driver 7x
- ▶ z196 and z114 with Driver 86 or Driver 93

4. The Optical Power Measurement panel at z10 with Driver 7x and at z196 and z114 with Driver 86 or Driver 93 are shown in Figure E-18. The optical signal power value is presented with mW (milliwatt) at z10 while it is presented with dBm at z196 and z114.

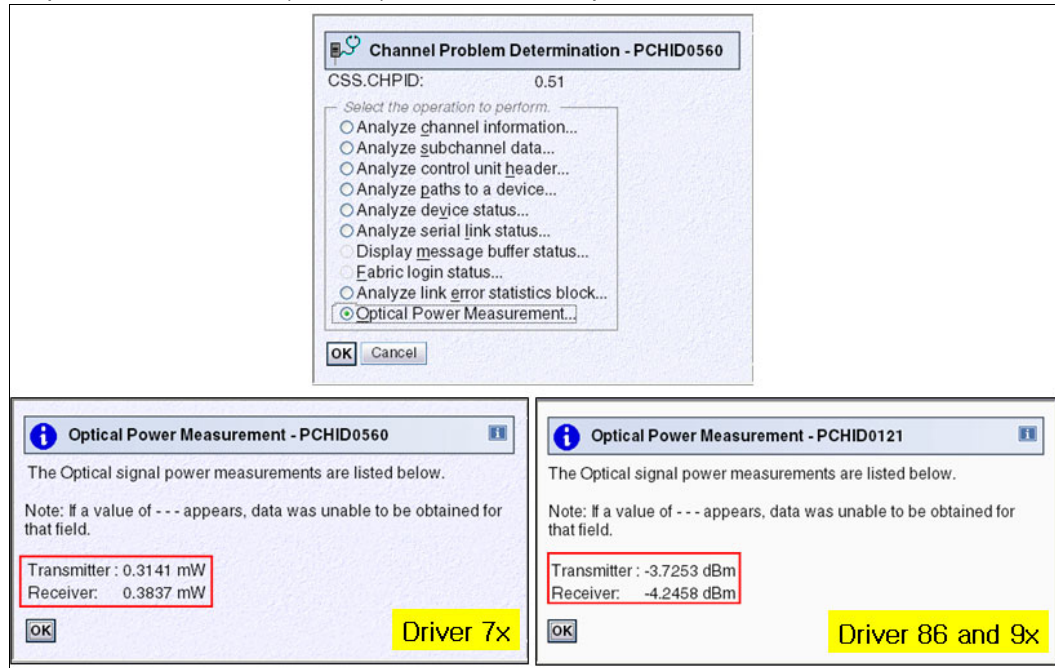


Figure E-18 Optical Power Measurement panel at Driver 7x,86,and 93

5. Table E-7 Minimum and maximum acceptable power specifications lists the acceptable optical power specifications for the FICON. The reference information is based on the Minimum and maximum acceptable power specifications in dBm at the book, *Maintenance Information for Fiber Optic Links*, SY27-2597-15 at Resource Link:

<http://halcontent.endicott.ibm.com/servers/resource1ink/svc03100.nsf/?OpenDatabase>

**Note:** The formula to convert from dBm to mW:  $\text{dBm} = 10 \log(k) \text{miliwatts}$  where k is optical power in miliwatt. The dBm value at the book, *Maintenance Information for Fiber Optic Links*(SY27-2597-15) is converted into mW with the formula in Table E-7 on page 451.

Table E-7 Minimum and maximum acceptable power specifications

| Link type                                       | TX Min                 | TX Max               | RX Min                | RX Max               |
|---|------------------------|----------------------|-----------------------|----------------------|
| Multimode FICON LX with mode conditioning patch | -8.5 dBm<br>(0.141 mW) | -4 dBm<br>(0.398 mW) | -22 dBm<br>(0.006 mW) | -3 dBm<br>(0.501 mW) |

| Link type                          | TX Min                  | TX Max               | RX Min                  | RX Max               |
|------------------------------------|-------------------------|----------------------|-------------------------|----------------------|
| Single-mode FICON LX 1 Gb          | -9.5 dBm<br>(0.112 mW)  | -3 dBm<br>(0.501 mW) | -20 dBm<br>(0.01 mW)    | -3 dB<br>(0.501 mW)  |
| Single-mode FICON LX 2 Gb          | -11.7 dBm<br>(0.068 mW) | -3 dBm<br>(0.501 mW) | -20 dBm<br>(0.01 mW)    | -3 dBm<br>(0.501 mW) |
| Single-mode FICON LX 4 Gb<br>10 km | -8.4 dBm<br>(0.145 mW)  | -1 dBm<br>(0.794 mW) | -16 dBm<br>(0.025 mW)   | -1 dBm<br>(0.794 mW) |
| Single-mode FICON LX 4 Gb 4 km     | -11.2 dBm<br>(0.076 mW) | -1 dBm<br>(0.794 mW) | -16 dBm<br>(0.025 mW)   | -1 dBm<br>(0.794 mW) |
| Single-mode FICON LX 8 Gb          | -8.4 dBm<br>(0.145 mW)  | -1 dBm<br>(0.794 mW) | -13.5 dBm<br>(0.045 mW) | -1 dBm<br>(0.794 mW) |
| Multimode FICON SX 1 Gb            | -10 dBm<br>(0.1 mW)     | -1 dBm<br>(0.794 mW) | -16 dBm<br>(0.025 mW)   | 0 dBm<br>(1 mW)      |
| Multimode FICON SX 2 Gb            | -10 dBm<br>(0.1 mW)     | -1 dBm<br>(0.794 mW) | -14 dBm<br>(0.04 mW)    | 0 dBm<br>(1 mW)      |
| Multimode FICON SX 4 Gb            | -9 dBm<br>(0.126 mW)    | -1 dBm<br>(0.794 mW) | -13 dBm<br>(0.05 mW)    | 0 dBm<br>(1 mW)      |
| Multimode FICON SX 8 Gb            | -8.2 dBm<br>(0.151 mW)  | -1 dBm<br>(0.794 mW) | -9.5 dBm<br>(0.112 mW)  | 0 dBm<br>(1 mW)      |

6. When a channel link is connected through a FICON Switch, You can display the Optical Power level at Switch side using a Command through CLI (Command Line Interface) or GUI (Graphic User Interface) based tool (for instance IBM Network Advisor). The CLI command and GUI tool are specific for each Switch Manufacturer. Figure E-19 on page 453 shows Brocade FICON Switch Optical Power level as a sample. Now the link loss budget can be calculated without the need for measuring the link directly by installing a power meter. The FICON Switch side value is depending on the switch manufacturer, type and model and version level of the switch firmware:

- To calculate the loss on any given link, the values from both sides of the link are needed. When both sets of values are known, the loss for the link can be calculated:

\* (TX value of CPU side) minus (RX value of FICON Switch side)

\* (TX value of FICON Switch side) minus (RX value of CPU side)

- The result here is the link loss and it is depending on Card Type, SFP type, Mode, and distance. Refer to Table 4-1 on page 85 for the link loss budget.

## Brocade FICON Switch SFP Power level sample

### Network Advisor

Network Advisor interface showing SFP power levels for Port 0. The 'Advanced' tab is selected, showing TX Power at 851.8 uWatts (-1.859 dBm) and RX Power at 645.6 uWatts (-1.9 dBm).

### CLI command: sfpshow slot#/port#

```

FID001_SAN768B_2: FID1:admin> sfpshow 1/0
Identifier: 3 SFP
Connector: 7 LC
Transceiver: 5401001200000000 2.4.8_Gbps SM lw Long_dist
Encoding: 1 8B10B
Baud Rate: 85 (units 100 megabaud)
Length 9a: 10 (units km)
Length 9b: 100 (units 100 meters)
Length 50a: 0 (units 10 meters)
Length 62.5a: 0 (units 10 meters)
Length Cu: 0 (units 1 meter)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000027-01
Vendor Rev: A
Wavelength: 1310 (units nm)
Options: 001a Loss_of_Sig,Tx_Fault,Tx_Disable
BR Max: 0
BR Min: 0
Serial No: UDA111041032041
Date Code: 110211
DD Type: 0x68
Enh Options: 0xf0
Status/Ctrl: 0xb2
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0

Temperature: 43 Centigrade
Current: 31.552 mAmps
Voltage: 3251.2 mVolts

Alarm          Warn
low  high     low  high
-40  90        -10  85
2.000  95.000    2.000  80.000
2700.0 3800.0    2970.0 3630.0

RX Power: -1.9 dBm (641.1uW)  4.4 uW  2240.0 uW  11.2 uW  1122.0 uW
TX Power: -1.9 dBm (651.8 uW)  27.0 uW  2240.0 uW  67.0 uW  1122.0 uW
                    
```

Figure E-19 Sample Optical Power Display at Brocade FICON Switch

## E.4 Displaying FICON IFCC and Other errors

You can see FICON-related error information in detail and some recovery action to fix a FICON failure at the IFCC and Other panel.

- ▶ When you select a PCHID, all different types of FICON errors are displayed for the PCHID, as shown at Figure E-20 on page 454. Typical FICON error types are:
  - FICON IFCC detected
  - FICON ELS (Extended Link services) error detected
  - FICON timeout
  - FICON External error
  - Channel Hardware failure
  - PPE (Purge Path Extended) Log

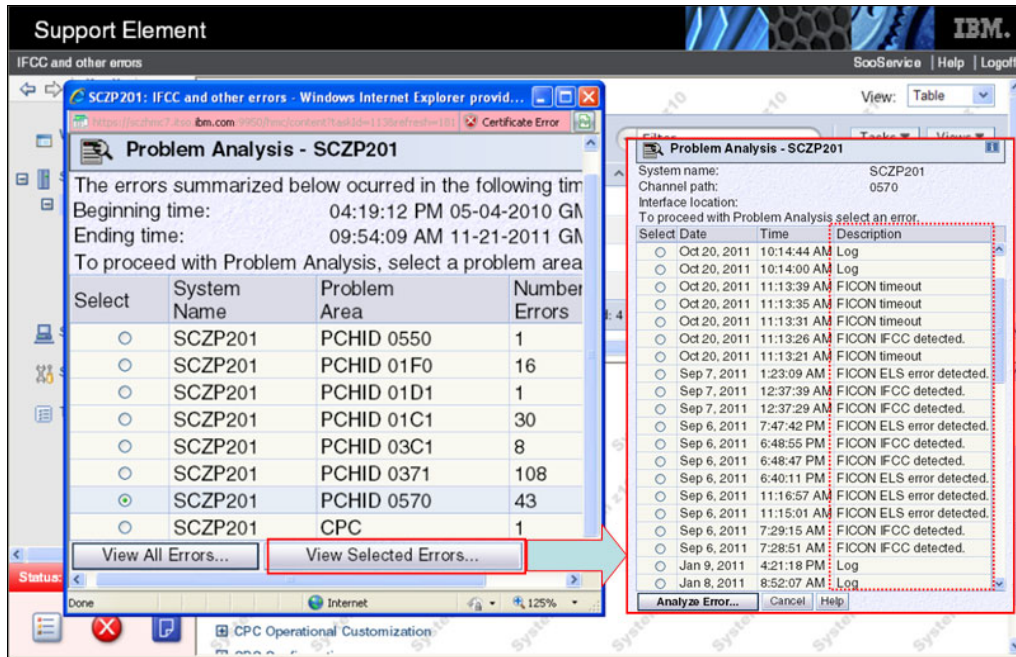


Figure E-20 IFCC and Other panel

## E.4.1 Displaying the IFCC and Other panel

To open the IFCC and Other panel:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. On the SE Navigation pane, click **System Management**, and click a CPC to work with.
3. Click the **IFCC and Other** at Service task list at Tasks pad.

Figure E-20 shows the **IFCC and Other panel**.

The Problem Analysis window displays the dates and times that errors occurred on the selected PCHID. Select a FICON PCHID associated with the error (usually the one nearest the time of the failure), and click **View Selected Errors**.

## E.4.2 FICON IFCC detected error

When you want to examine a FICON IFCC detected error, follow these steps to open the FICON IFCC detected error, as shown in Figure E-21 on page 455:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the IFCC and Other panel (see Displaying the IFCC and Other panel on page 454).
3. Click the **FICON IFCC detected error** you want to examine, and click **Analyze error**.
4. The Problem Analysis window gives a description of the causes of the error and a possible corrective action. To see the trace taken at the time of the error, click **I/O Trace**, or **More information** depending on the panel.
5. The Input/Output (I/O) Trace Results window provides a trace of the interface at the time of the error, such as System Reference Code (SRC) for further analysis.

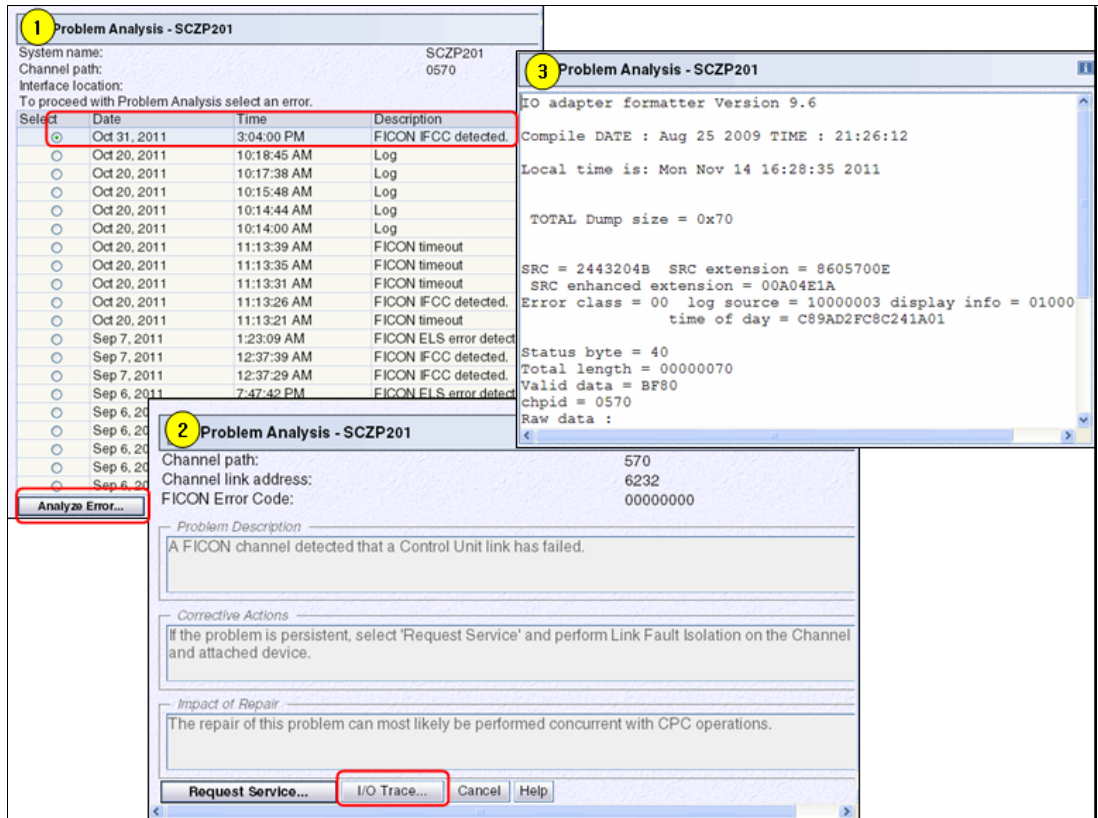


Figure E-21 FICON IFCC detected error panel

### E.4.3 FICON ELS (Extended Link Services) error detected error

When you want to examine a FICON ELS error detected error, follow these steps to open the FICON ELS error detected error shown in Figure E-22 on page 456 and Figure E-23 on page 457:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the IFCC and Other panel (see Displaying the IFCC and Other panel on page 454).
3. Click the **FICON ELS error detected error** you want to examine and click **Analyze error**.
4. The Problem Analysis window gives a description of the causes of the error and a possible corrective action. To see the trace taken at the time of the error, click **I/O Trace** or **More information** depending on the panel as shown at Figure E-22 on page 456.
5. The Input/Output (I/O) Trace Results window provides a trace of the interface at the time of the error, such as System Reference Code (SRC) for further analysis, as shown in Figure E-23 on page 457.
6. The Input/Output (I/O) Trace Results window also provides more informations regarding the ELS error:
  - Channel information: CHPID and Link address
  - ELS information: the failed ELS name and Return code
  - For a detailed description for more ELS information and the debugging of these records, refer to the ANSI standard architecture, which can be found at: <http://www.t11.org>

**1 Problem Analysis - SCZP201**

System name: SCZP201  
Channel path: 0570  
Interface location:  
To proceed with Problem Analysis select an error.

| Select                           | Date         | Time        | Description               |
|----------------------------------|--------------|-------------|---------------------------|
| <input type="radio"/>            | Oct 31, 2011 | 3:04:00 PM  | FICON IFCC detected.      |
| <input type="radio"/>            | Oct 20, 2011 | 10:18:45 AM | Log                       |
| <input type="radio"/>            | Oct 20, 2011 | 10:17:38 AM | Log                       |
| <input type="radio"/>            | Oct 20, 2011 | 10:15:48 AM | Log                       |
| <input type="radio"/>            | Oct 20, 2011 | 10:14:44 AM | Log                       |
| <input type="radio"/>            | Oct 20, 2011 | 10:14:00 AM | Log                       |
| <input type="radio"/>            | Oct 20, 2011 | 11:13:39 AM | FICON timeout             |
| <input type="radio"/>            | Oct 20, 2011 | 11:13:35 AM | FICON timeout             |
| <input type="radio"/>            | Oct 20, 2011 | 11:13:31 AM | FICON timeout             |
| <input type="radio"/>            | Oct 20, 2011 | 11:13:26 AM | FICON IFCC detected.      |
| <input checked="" type="radio"/> | Sep 7, 2011  | 1:23:09 AM  | FICON ELS error detected. |
| <input type="radio"/>            | Sep 7, 2011  | 12:37:39 AM | FICON IFCC detected.      |
| <input type="radio"/>            | Sep 7, 2011  | 12:37:39 AM | FICON IFCC detected.      |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |
| <input type="radio"/>            | Sep 6, 2011  |             |                           |

**2 Problem Analysis - SCZP201**

Channel path: 570  
FICON Error Code: 00000000

*Problem Description*

A FC or FCV channel has detected an error in an extended link service request or response that is unrelated to link initialization.

*Corrective Actions*

If the problem is persistent, select 'Request Service' and Contact next level of support.

*Impact of Repair*

The repair of this problem can most likely be performed concurrent with CPC operations.

Figure E-22 FICON ELS error detected error panel (1/2)



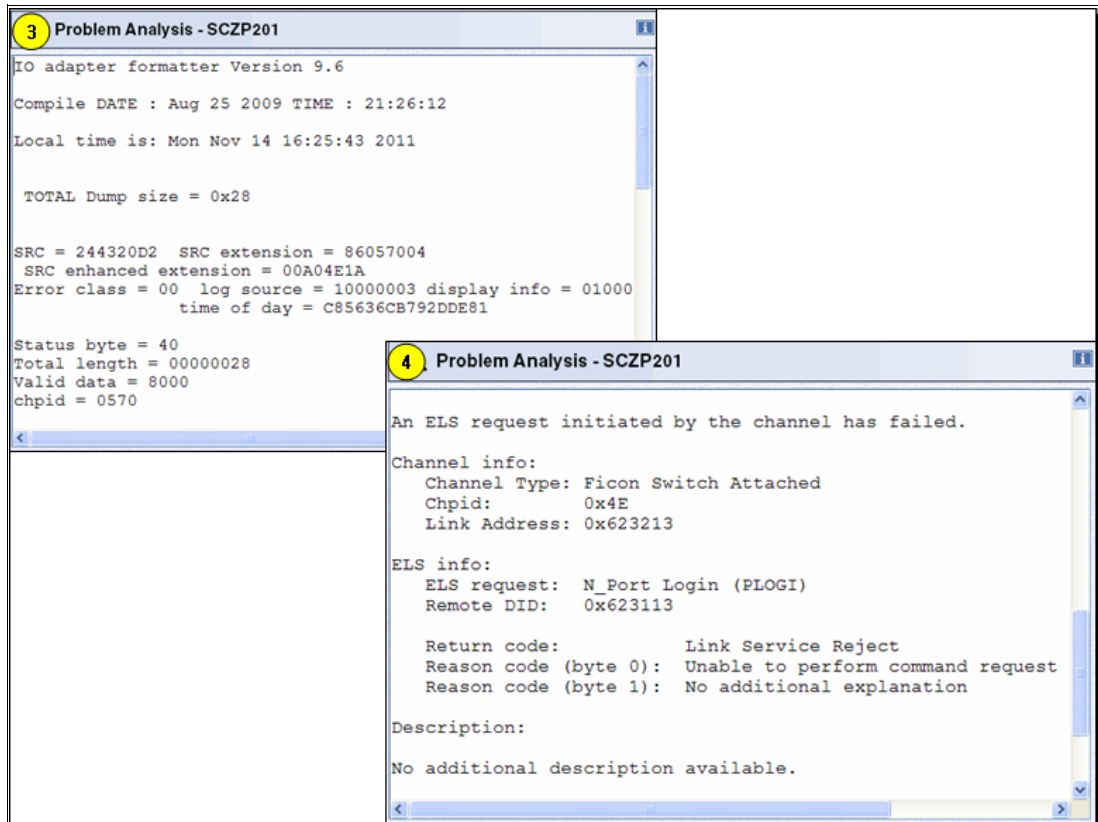


Figure E-23 FICON ELS error detected error panel(2/2)

#### E.4.4 FICON timeout error

When you want to examine a FICON timeout error, follow these steps to open the FICON timeout error shown at the Figure E-24 on page 458 and Figure E-25 on page 459.

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the **IFCC and Other** panel (see Displaying the IFCC and Other panel on page 454).
3. Click the **FICON timeout error** you want to examine, and click **Analyze error**.
4. The Problem Analysis window gives an error code(90FCFEFE), which is a description of the causes of the error and a possible corrective action, as shown at Figure E-24 on page 458. To see the trace taken at the time of the error, click **I/O Trace** or **More information** depending on the panel, as shown at Figure E-25 on page 459.

The Input/Output (I/O) Trace Results window provides a IU (Information Unit) trace of the interface at the time of the error for further analysis, such as:

- IUI (Information Unit ID): IU type and IU function identified:
  - Type: Data/CMD/Status/Control/CMD-Data/Link Control
  - AS (Address Specific): IU is associated with specific device
  - ES (Supplemental Status): Meaning depends on IU Type
- DHF (Device Header Flag): To invoke SB protocols to control the I/O operation:
  - E (End): Meaning depends on IU Type
  - EC (End Connection): Connection end

- RDY: Ready
  - CH (Chaining): To signal Data or CMD chaining
  - EE (Early End): Meaning depends on IU Type
  - CNP(CRC Not Provided): To indicate the value contained in the CRC field is not valid
- DIB (Device Information Block):
    - B0-B3(Byte 0-3): DIB specific
    - B7(Byte 7): Command flags for CMD/ CMD-Data DIB
    - IU#: Number of IU count in a Exchange
    - Data BC (Data Byte Count): DIB data byte count and DIB specific
  - CHI (Channel Image): Image ID
  - CUI (Control Unit Image): CUADD in IOCDs
  - UA (Unit Address): Device Unit Address
  - LA (Link Address): Destination Link Address
  - CMD#: Exchange CCW CMD number
  - CCW#: Unique CCW number for a IU
  - Token: Exchange Token
  - Dev#: Actual Device Number in IOCDs
- More details at: FC-SB4 documentb

<http://www.t11.org>

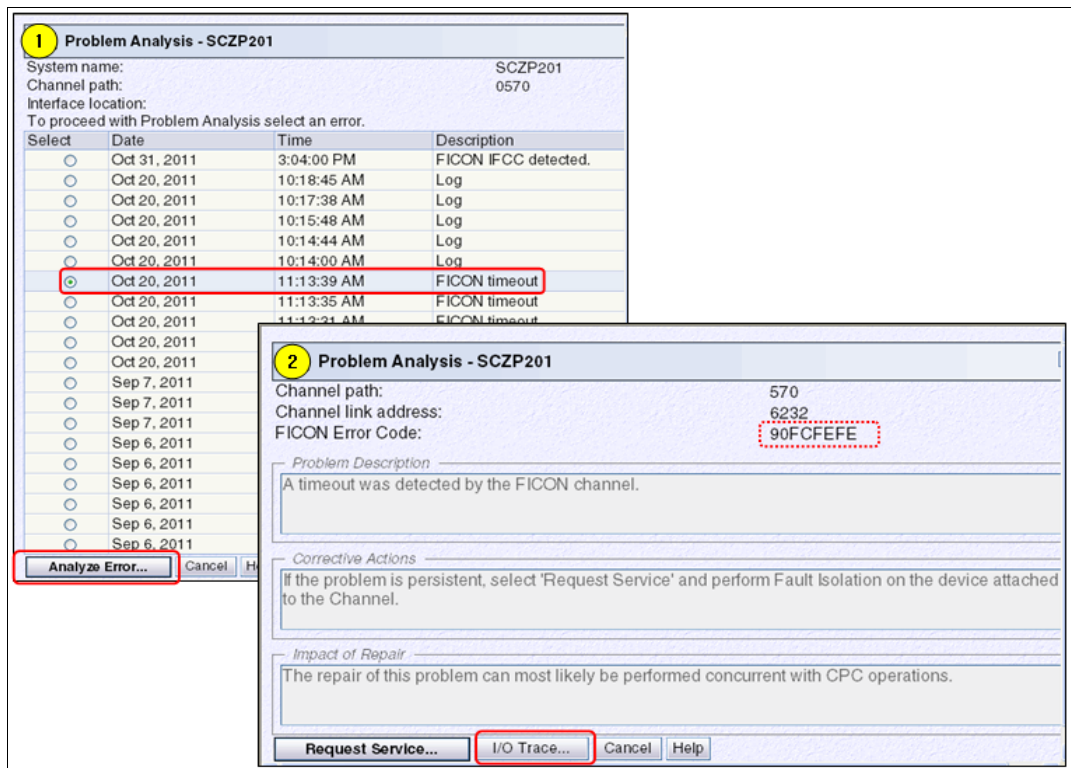


Figure E-24 FICON timeout error panel (1/2)

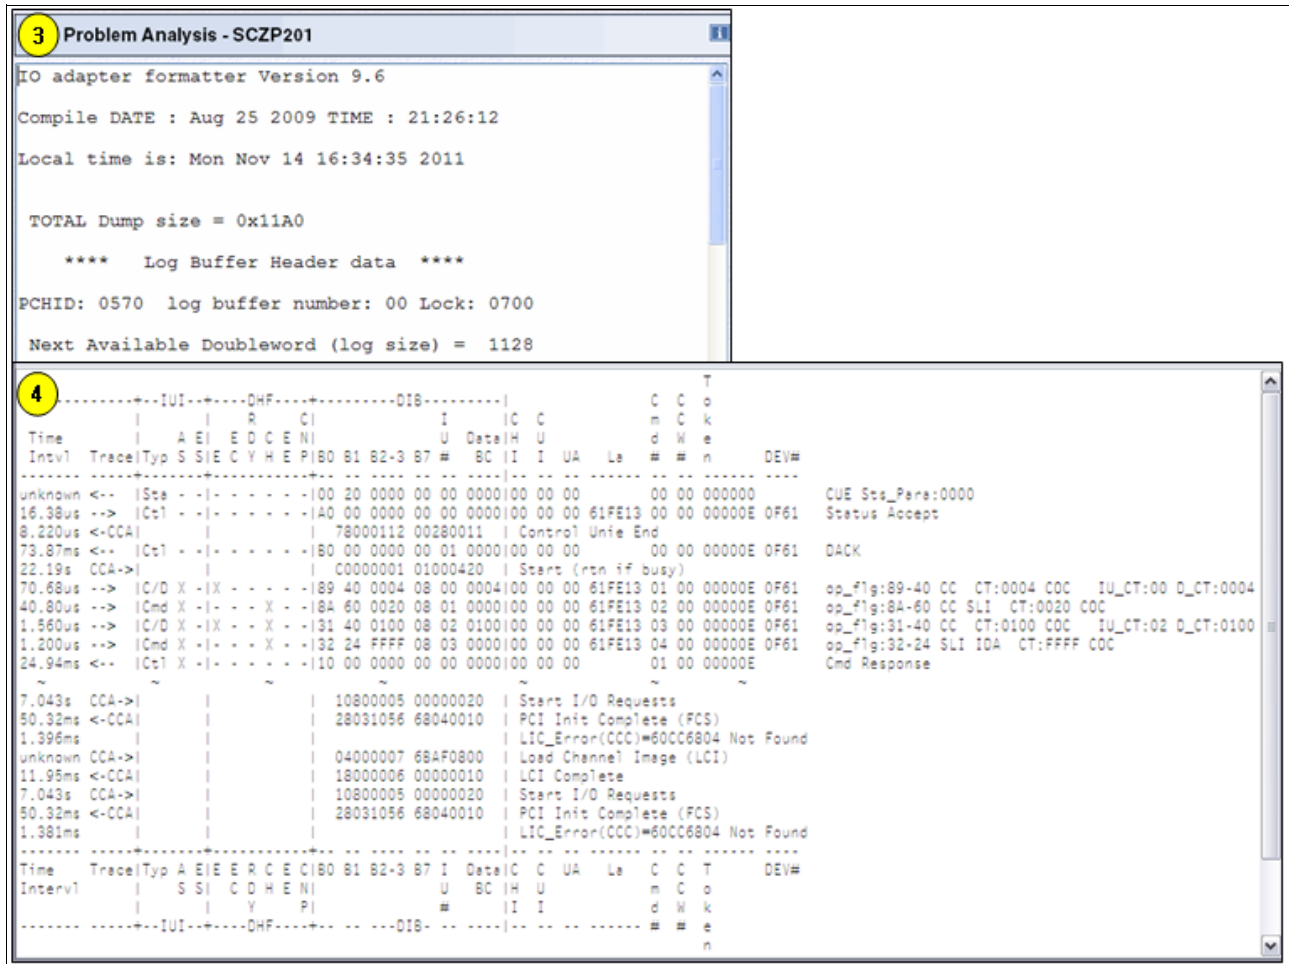


Figure E-25 FICON timeout error panel(2/2)

### E.4.5 FICON External error

When you want to examine a FICON External error, follow these steps to open the FICON External error shown at the Figure E-26 on page 460 and Figure E-27 on page 460:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the **IFCC and Other** panel (see Displaying the IFCC and Other panel on page 454).
3. Click the **FICON timeout error** you want to examine, and click **Analyze error**.
4. The Problem Analysis window gives an error code(90FCAB00),a description of the causes of the error and a possible corrective action as shown at Figure E-26 on page 460. To see the trace taken at the time of the error, click **I/O Trace** or **More information** depending on the panel, as shown at Figure E-27 on page 460.

The Input/Output (I/O) Trace Results window provides a IU (Information Unit) trace of the interface at the time of the error for further analysis same as the one for FICON timeout error. Refer to the details of FICON timeout error trace at FICON timeout error on page 457.

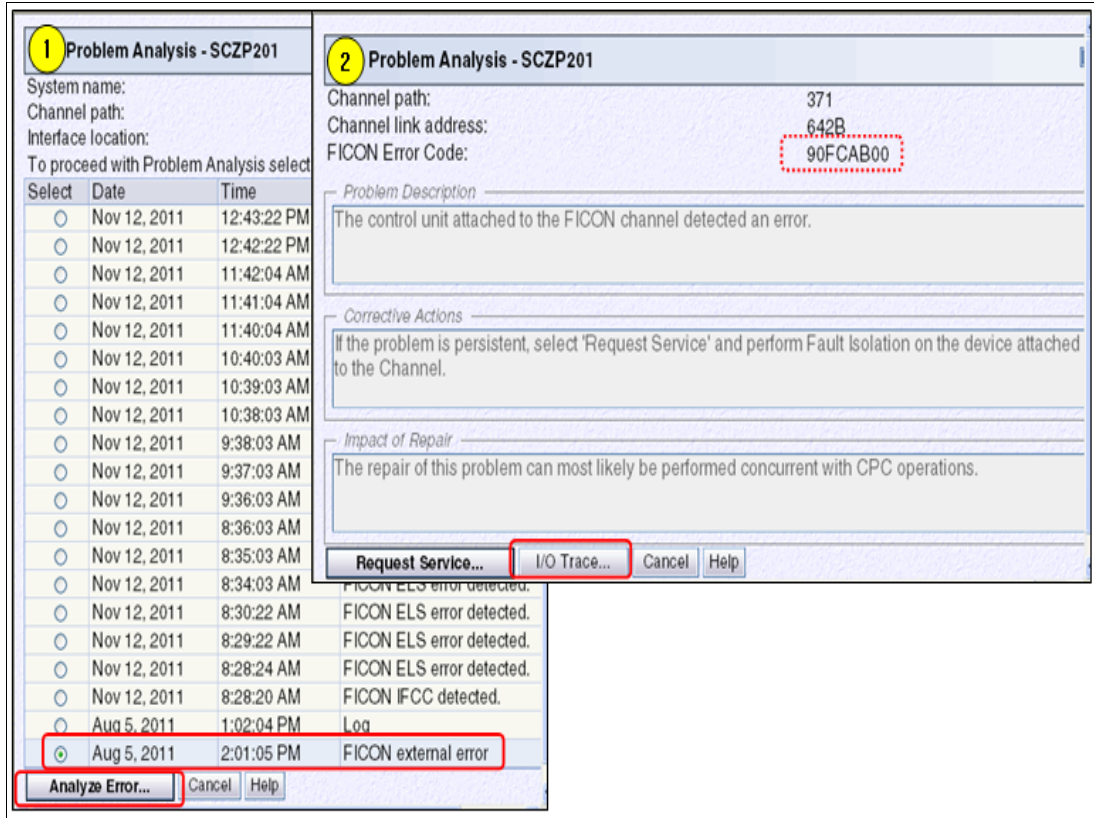


Figure E-26 FICON External error panel(1/2)

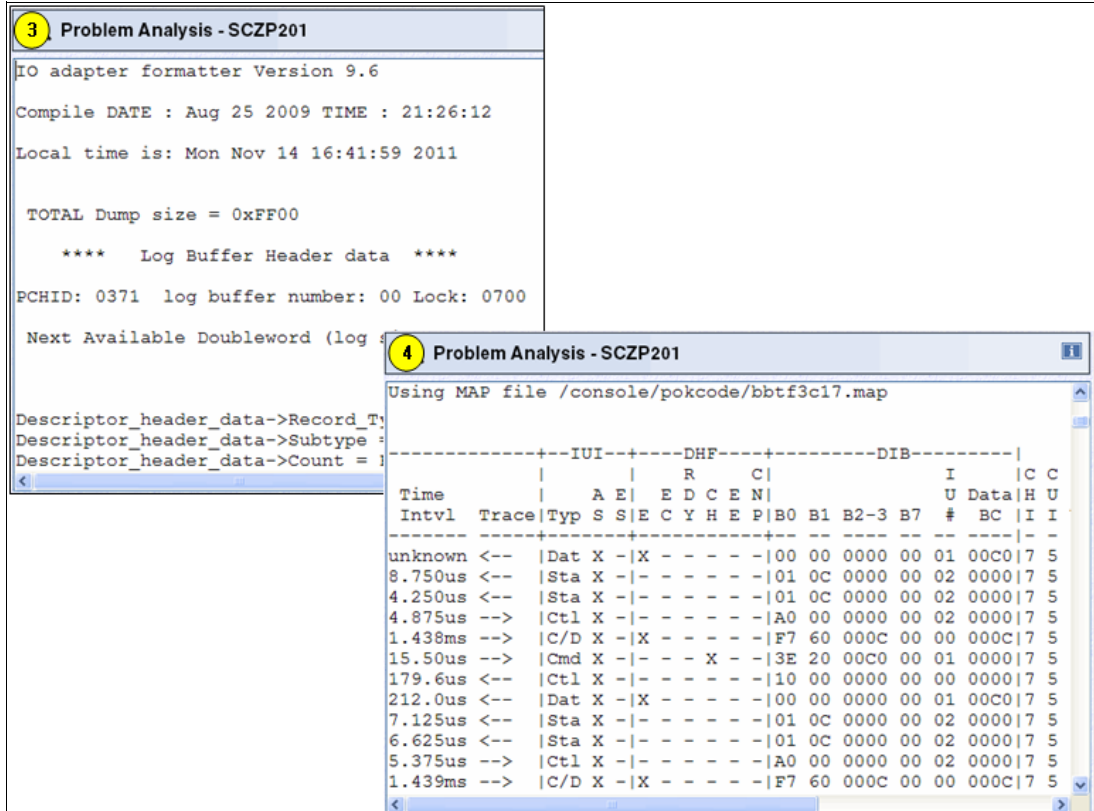


Figure E-27 FICON External error panel(2/2)

## E.4.6 Channel Hardware Failure

When you want to examine a Channel Hardware Failure, follow these steps to open the Channel Hardware Failure shown in the Figure E-28:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the **IFCC and Other** panel (see Displaying the IFCC and Other panel on page 454).
3. Check a channel PCHID with its Hardware location (for instance, Z01BLG04J.00). Click the **Channel Hardware Failure** you want to examine, and click **Analyze error**.
4. The Problem Analysis window gives an error code (for instance, 90FCFEFE), which is a description of the causes of the error and impact of repair. To see the trace taken at the time of the error, click **I/O Trace** or **More information** depending on the panel, as shown at Figure E-28.
5. You can see more information to get the support of an IBM Service Representative, such as SRC information (for instance, 2243200C).

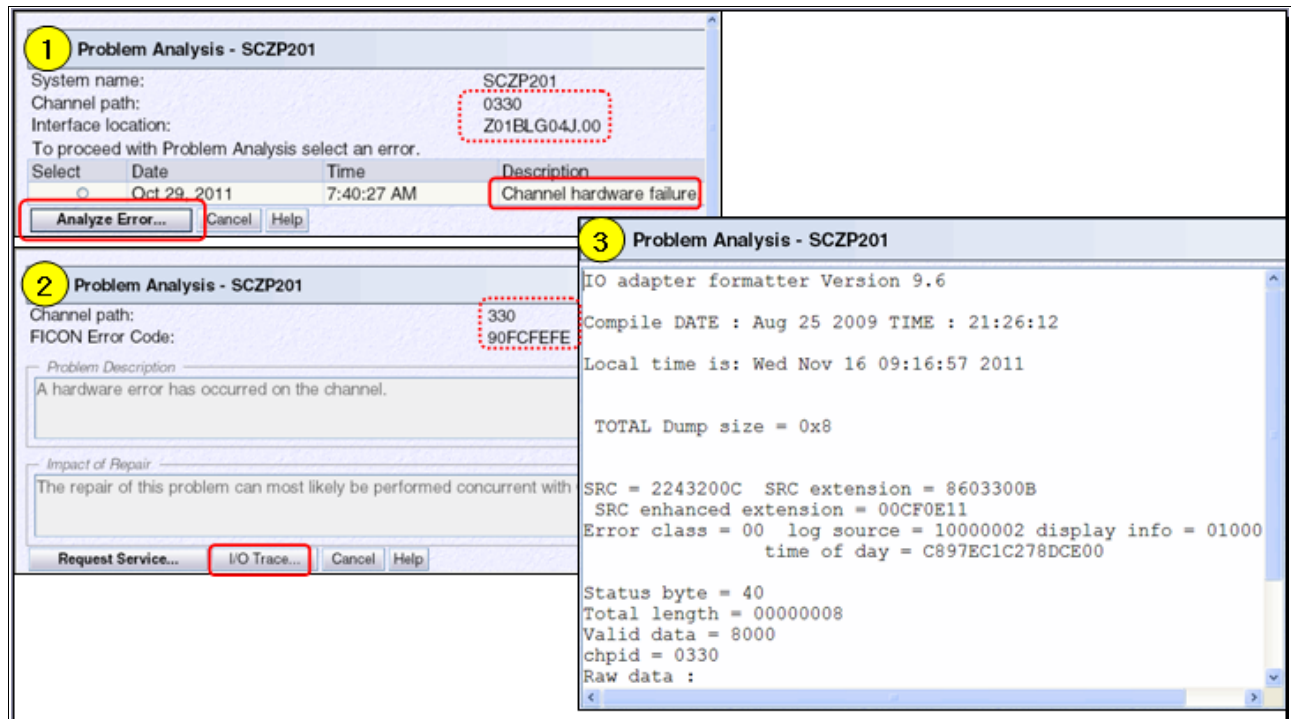


Figure E-28 Channel Hardware Failure panel

## E.4.7 Purge Path Extended Log

When you want to examine a PPE log, follow these steps to open the PPE log shown in the Figure E-29 on page 462:

1. Log on to the HMC and SE in service mode (see Logging on to the HMC and SE on page 427).
2. Display the IFCC and Other panel (see Displaying the IFCC and Other panel on page 454).
3. Click the description type, Log you want to examine, and click **Analyze error**.

- The Problem Analysis window gives a description of the error as Purge Path Extended Data. To see the trace taken at the time of the error, click **I/O Trace** as shown in Figure E-29.

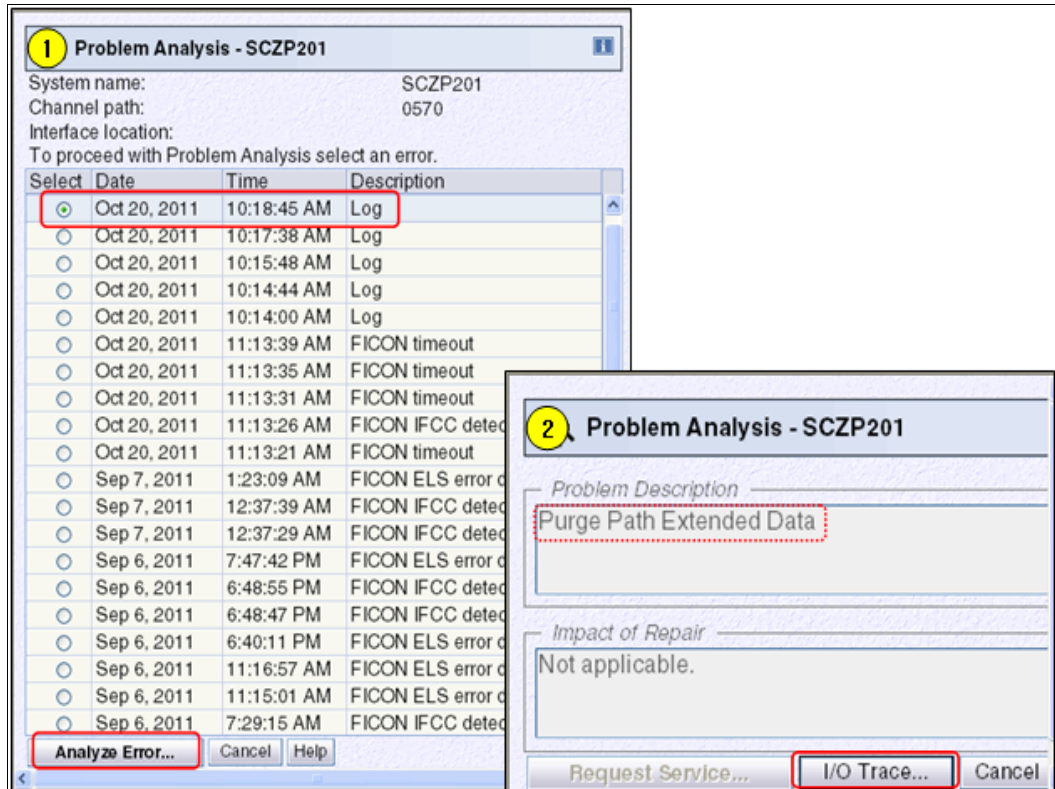


Figure E-29 Purge Path Extended Log panel

Figure E-30 and Figure E-31 on page 464 show the details of a PPE trace sample.

```

Purge Path Extended Log

PPE Header_____
Lock_____0800   Log Data Count_____0210
Log Code_____F8   Path Mask_____01
Buffer Number_____00
Validity Flags_____38   PCHID_____0570

20 CHANNEL NPORT Only Valid data in this Log
10 Channel Fport
08 Channel ModDep
04 CU Nport
02 CU Fport
01 CU ModDep

Link Address_____016D CU Logical Address_____02
CSSID_IID_____0C   Device Address_____B204
The CU Link Address in this case would be Switch 01 Port 6D

Channel LESB_____ <--CHPID N PORT
Link Failure Count_____00000000 Loss of Sync Count_____00000005
Loss of Sgnl Count_____00000000 Prim Seq Err Count_____00000000
Inv Xmit Wrđ Count_____00000010 Inv CRC Err Count_____00000000

Channels FPort LESB_____ <--SWITCH F PORT
Link Failure Count_____0000003D Loss of Sync Count_____0000001D
Loss of Sgnl Count_____00000024 Prim Seq Err Count_____00000000
Inv Xmit Wrđ Count_____0000009C Inv CRC Err Count_____00000000
Channel Model Dependent___
Architected Error Code__00 (This would be from FC-SB-4 Doc.)
(00 = No Error was recognized)
Reserved_____00 Reserved_____0000
Caller ID_____90 Reserved_____00
Local Error Code_____FEFE (The Channel reported IFCC, Timeout 90FCFEFE)
CU Link Addr_____00016D00 Channel Link Addr_____00000103
The CU attachment is on Switch 01 Port 6D
The CH attachment is on Switch 01 Port 03
FPort Node WWN_____
0000 : 21000060 DF224DF6
FPort Port WWN_____
0000 : 21000060 DF224DF6
NPort Port WWN_____
0000 : 50050764 01007AC8
NPort Node WWN_____
0000 : 50050764 00C6943D

```

Figure E-30 PPE Trace sample(1/2)

```

CU RNID data
Function_____02000000 Node ID Data Format_____18
Common Node ID Length___10 Reserved_____00
Specific Node ID Length__20
NPORT Name_____
0000 : 50050763 03104015
Node Name_____
0000 : 50050763 03FFC015
Flags_____00
Node Parameters_____
0000 : 200100
Type Number_____
0000 : 30303231 3037 (002107)<--CU Machine Type
Model Number_____
0000 : 393231 (921)<--CU Machine Serial Number
Manufacturer_____
0000 : 49424D<--IBM (Brocade: 4D4344)
Plant of Manufacture_____
0000 : 3133 (13)
Sequence Number_____
0000 : 30303030 30303030 31323331 (000000001231) Serial #
Tag_____
0000 : 0201 (CU I/F Slot)
CU Logout_____
CU NPort LESB_____ <--CU N PORT
Link Failure Count_____00000000 Loss of Sync Count_____00000000
Loss of Sgnl Count_____00000000 Prim Seq Err Count_____00000000
Inv Xmit Wrđ Count_____00000000 Inv CRC Err Count_____00000000
CU FPort LESB_____ <--SWITCH F PORT
Link Failure Count_____
00000000 Loss of Sync Count_____00000000
Loss of Sgnl Count_____00000000 Prim Seq Err Count_____00000000
Inv Xmit Wrđ Count_____00000000 Inv CRC Err Count_____00000000

```

Figure E-31 PPE Trace sample(2/2)

Reference information for LESB analysis is in the FC-PH Doc. Link failure counters:

- PPE does not reset the LESB counters. So multiple events need to be reviewed. Otherwise you might be looking at OLD ERROR INCIDENTS.
- Four types of Link failures are recorded in individual counters in LESB. The Link failure counters are:
  - Link Failure Count (miscellaneous link errors)
  - Loss of Synchronization Count (confirmed and persistent sync. loss)
  - Loss of Signal Count
  - Primitive Sequence Protocol Error Count
- Invalid Transmission Word: The Invalid Transmission Word Counter increments, once for every Invalid Transmission Word received, except during the following conditions:
  - No Transmission Word errors are counted if the receiver is in the Loss of Synchronization state.
  - No Transmission Word errors are counted if the Port is in the OLS (Offline Signal) Receive state (OL2) or the Wait for OLS state (OL3).



- Invalid CRC Count:
- The Invalid CRC Count increments, once for every received frame which meets one of the following conditions:
  - The Port is in the Active State (AC) and the received frames CRC is in error and the frame is either missing an EOF (End Of Frame) delimiter or the EOF delimiter is an EOFn, EOFt, or EOFdt
  - The Port is in the Active State (AC) and the received frames CRC is in error.

The Architected Error Code is in FC-SB4 Doc.

Error Codes for the Purge-Path-Response IU:

- 00 No errors were recognized
- 01 SB-4 protocol timeout
- 02 SB-4 link failure
- 03 Logical-path-timeout error
- 04 SB-4 offline condition
- 05 FC-FS-3 link failure
- 06 SB-4 length error
- 07 LRC error
- 08 SB-4 CRC error
- 09 IU count error
- 0A SB-4 link-level protocol error
- 0B SB-4 device-level protocol error
- 0C Receive ABTS <---- in this sample trace case
- 0D Reserved
- 0E Abnormal termination of exchange
- 0F Logical path not established
- 10 Test initialization result error
- 11-FF Reserved

Error Codes for the Purge-Path IU:

- 00 Error-code transfer not supported
- 01 SB-4 protocol timeout
- 02 SB-4 link failure
- 03 Reserved
- 04 SB-4 offline condition
- 05 FC-FS-3 link failure
- 06 SB-4 length error
- 07 LRC error
- 08 SB-4 CRC error
- 09 IU count error
- 0A SB-4 link-level protocol error
- 0B SB-4 device-level protocol error
- 0C Receive ABTS
- 0D Cancel function timeout
- 0E Abnormal termination of exchange
- 0F Host storage error
- 10 Software termination of exchange due to halt request
- 11 Software termination of exchange due to clear request
- 12 Interrogate operation error
- 13 Transport operation error
- 14 Transport error
- 15 REC error
- 16-FF Reserved

## E.5 Displaying the Fibre Channel Analyzer panel

The Fibre Channel Analyzer task is a new HMC task introduced on Driver 79 and enhancements to assist with FICON Channel Problem Determination. The function is controlled when you enable options in the HMC Customize Console Services task located in the HMC Console tasks. When Fibre Channel Analyzer is enabled on the HMC, all SEs defined to the HMC will report fiber events to the HMC.

To open the Fibre Channel Analyzer panel:

1. Log on to the HMC in sysprog or service mode (see Logging on to the HMC and SE on page 427).

To verify that the Fibre Channel Analysis function enabled:

2. On the HMC Navigation pane, click **HMC Management**.
3. On the HMC Work pane, click **Customize Console Services** at the Configuration task list.
4. Enable the **Fibre Channel Analysis** function, and click **OK**.

To display the Fibre Channel Analyzer pane:

On the HMC Navigation pane, click **Service Management**.

5. On the HMC Work pane, click **Fibre Channel Analyzer** at the Service Management task list.

The HMC will store the records in a component log on the HMC and display the records with the Fibre Channel Analyzer task on the HMC. When a threshold of errors has been exceeded on the HMC, a Fibre Channel Hardware Message will be posted on the HMC. The HMC will display a blue FICON ICON in the Work pane of the SE, the blue Fibre Channel Error ICON is used to open Hardware Message task as shown at Figure E-32 on page 467. Scroll down the error description at the Hardware Message task and you can see “FICON Analyzer Application to view the collected errors”.

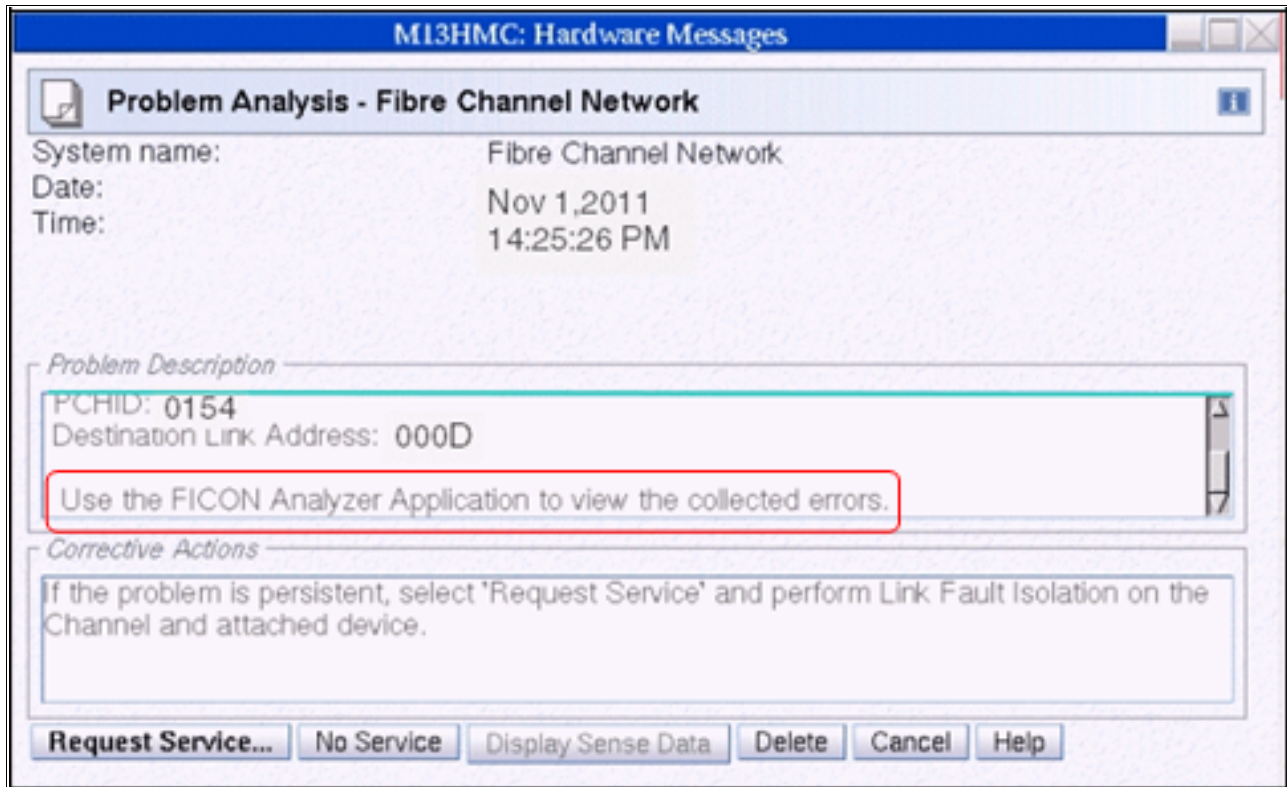


Figure E-32 Fibre event record at the HMC Hardware message icon

6. Double-click **Fibre Channel Analyzer** task in the Console Actions. Click the associated CPC and PCHID with the FICON error event you want to examine. You can see more detailed error information, as shown in Figure E-33 on page 468 and Figure E-34 on page 468 such as Channel and Device information and Reason code. The Reason code is displayed only in service mode logon.

**Fibre Channel Error Summary**

Select an entry to display all detailed information.

Fibre Channel Analyzer Display Panel

--- Select Action ---

| Select                           | System   | PCHID | Source Link Address | Destination Link Address | CHPID | Channel Type | Error Count |
|----------------------------------|----------|-------|---------------------|--------------------------|-------|--------------|-------------|
| <input checked="" type="radio"/> | P007E296 | 0154  | 000C                | 000D                     | 8A    | FICON        | 5           |
| <input type="radio"/>            | P007E296 | 015C  |                     |                          | 8E    | FICON        | 2           |
| <input type="radio"/>            | P007E296 | 0154  |                     |                          | 8A    | FICON        | 2           |

Total: 3 Filtered: 3 Selected: 1

Error Summary Output

```

1F0E00008ECF67605F 11/01/2011 14:32:56.564 IFCC on PCHID: 0154
* Device Address: 7A40 * Control Unit Id: 00 * Image Id: 2 * Channel Subsystem Id: 0
* Device MIMS: 3868-003-07E29601558B
-----
1F0E000090CD67602C 11/01/2011 14:32:55.487 IFCC - RC: 90FCFEFE on PCHID: 0154
* Device Address: 7A6E * Control Unit Id: 00 * Image Id: 2 * Channel Subsystem Id: 0
* Device MIMS: 3868-003-07E29601558B
-----
1F0E0000B2F360602D 11/01/2011 14:25:26.490 IFCC - RC: 90FCFEFE on PCHID: 0154
* Device Address: 7A07 * Control Unit Id: 00 * Image Id: 1 * Channel Subsystem Id: 1
* Device MIMS: 3868-003-07E29601558B

```

Figure E-33 Fibre Channel Analyzer panel(1/2) in sysprog mode

**Fibre Channel Error Summary**

Select an entry to display all detailed information.

Fibre Channel Analyzer Display Panel

--- Select Action ---

| Select                           | System   | PCHID | Source Link Address | Destination Link Address | CHPID | Channel Type | Error Count |
|----------------------------------|----------|-------|---------------------|--------------------------|-------|--------------|-------------|
| <input checked="" type="radio"/> | P007E296 | 0154  | 000C                | 000D                     | 8A    | FICON        | 5           |
| <input type="radio"/>            | P007E296 | 015C  |                     |                          | 8E    | FICON        | 2           |
| <input type="radio"/>            | P007E296 | 0154  |                     |                          | 8A    | FICON        | 2           |

Total: 3 Filtered: 3 Selected: 1

Error Summary Output

```

1F0E00008ECF67605F 11/01/2011 14:32:56.564 IFCC - RC: 90FCFEFE on PCHID: 0154
* Device Address: 7A40 * Control Unit Id: 00 * Image Id: 2 * Channel Subsystem Id: 0
* Device MIMS: 3868-003-07E29601558B
-----
1F0E000090CD67602C 11/01/2011 14:32:55.487 IFCC - RC: 90FCFEFE on PCHID: 0154
* Device Address: 7A6E * Control Unit Id: 00 * Image Id: 2 * Channel Subsystem Id: 0
* Device MIMS: 3868-003-07E29601558B
-----
1F0E0000B2F360602D 11/01/2011 14:25:26.490 IFCC - RC: 90FCFEFE on PCHID: 0154
* Device Address: 7A07 * Control Unit Id: 00 * Image Id: 1 * Channel Subsystem Id: 1
* Device MIMS: 3868-003-07E29601558B

```

Figure E-34 Fibre Channel Analyzer panel(2/2) in service mode

## E.6 Displaying the System I/O Configuration Analyzer panel

The System I/O Configuration Analyzer (SIOC analyzer) is a new HMC/SE-based tool that allows the system administrator to access the information from many sources in one place. This makes it much easier to manage I/O configurations, particularly across multiple CECs.

Through this function, you can see, PCHID, Control units, Link Load, and FC Node ID of current active IOCDs. Also, you can download Comma Separated Value files for each information for further analysis. You can find this menu under CPC Configuration of SE or Configuration of HMC.

To open the SIOC panel, log on to the HMC or SE in sysprog or service mode (see Logging on to the HMC and SE on page 427).

To display the SIOC panel at HMC:

1. On the HMC Navigation pane, click **System Management** and a **CPC** to work with.
2. On the HMC Tasks pad, click an SIOC task at the Configuration task list.

To display the SIOC panel at SE:

1. On the SE Navigation pane, click **System Management**, and click the **CPC**.
2. On the SE Tasks pad, click **SIOC** task at the CPC Configuration task list.

Different functions are at the top menu bar, as shown at Figure E-35:

When you display the SIOC Analyzer, the default screen is PCHID Control Unit View option.

| PCHID | Control units | 2    | 3    | 4    |
|-------|---------------|------|------|------|
| 0100  | 1             |      |      |      |
| 0101  | 1             |      |      |      |
| 0101  | 1             |      |      |      |
| 0101  | 1             |      |      |      |
| 0110  | 1             |      |      |      |
| 0111  | 1             |      |      |      |
| 0118  | 1             |      |      |      |
| 0118  | 1             |      |      |      |
| 0119  | 1             |      |      |      |
| 0120  | 1             |      |      |      |
| 0121  | 14            | CA00 | CC00 | CE00 |
| 0121  | 14            | CA00 | CC00 | CE00 |
| 0121  | 14            | CA00 | CC00 | CE00 |

Figure E-35 SIOC (System I/O Configuration) Analyzer panel

In Figure E-35 on page 469:

- ▶ File menu: Using “Save Data to USB Flash Memory Drive” option, you can download the analyzed file to CSV (Comma Separated Value) file format.
- ▶ View menu: There are five different format at the view menu.you can display the I/O configuration information based on what you want to check as shown at Figure E-36:
  - PCHID Control Unit: This is default screen when SIOC Analyzer task performed. All of the defined Control Unit Numbers are listed with PCHID,CSS.CHPID,Chanel Type, Switch, and Number of Control Units.
  - PCHID Partition: Lpar names are listed with PCHID,CSS.CHPID,Chanel Type, Switch, and Number of Control Units.
  - Control Unit: Each of Control Unit numbers are listed with the attached CHPIDs.
  - Link Load: Each of the defined Link address are listed with the defined PCHIDs:
    - Link: The attached Switch link address.
    - Number: Total defined number of PCHIDs for the Link.
    - PCHID (Defined Count): Total defined number of a PCHID for the Link
  - Node ID: The attached Node Type and its Protocol are listed with the PCHID and Link address.
- ▶ Filter menu: You can filter the I/O Configuration by each type of the filter options(PCHID,CSS.CHPID,Switch,Partition,and Control Unit)
- ▶ Sort menu: You can sort the Link Load View information with any combination of Number and Link option.

The screenshot shows the SIOC Analyzer View menu panel with five view options selected. Red arrows point from each option to its corresponding data table:

- PCHID Control Unit:**

| PCHID | CSS.CHPID | Type | Switch | Number of Control Units | Control units | 2 |
|-------|-----------|------|--------|-------------------------|---------------|---|
| 0100  | 1.03      | OSD  |        | 1                       | 20A0          |   |
| 0101  | 0.0B      | OSM  |        | 1                       | 2360          |   |
- PCHID Partition:**

| PCHID | CSS.CHPID | Type | Switch | Number of Control Units | LPAR Name |     |     |
|-------|-----------|------|--------|-------------------------|-----------|-----|-----|
| 0100  | 1.03      | OSD  |        | 1                       | A11       | A13 | A16 |
| 0101  | 0.0B      | OSM  |        | 1                       | A01       | A02 | A03 |
- Control Unit:**

| Control Unit Number | Attached PCHIDs                |  |  |  |                     |  |  |
|---------------------|--------------------------------|--|--|--|---------------------|--|--|
| 0B00.               | 01D1.( *6311/*6311/*6311/----) |  |  |  |                     |  |  |
| 0B01.               | 0553.( *6411/*6411/*6411/----) |  |  |  |                     |  |  |
| 0B90.               | 0160.( *AEEC/----/----/----)   |  |  |  | 0164.(----/----/*AE |  |  |
- Link Load:**

| Link | Number | PCHID (Defined Count) |         |         |         |         |
|------|--------|-----------------------|---------|---------|---------|---------|
| 0000 | 41     | 0110(1)               | 0111(1) | 0118(2) | 0119(1) | 0190(1) |
| 000D | 66     | 0100(1)               | 0101(3) | 0120(1) | 0180(4) | 0181(1) |
- Node ID:**

| PCHID | Link | Validity | Node Type | Protocol |
|-------|------|----------|-----------|----------|
| 0100  | 0001 | Valid    | Device    | FC-SB-2  |
| 0110  | 0000 | Valid    | CPC       | FC-SB-2  |
| 0111  | 0000 | Valid    | CPC       | Reserved |

Figure E-36 SIOC Analyzer View menu panel



## Useful z/OS commands

We provide information pertaining to the FICON environment that can be obtained from various z/OS commands.

# Using z/OS commands for problem determination

Figure F-1 represents the FICON environment that we used to produce the output of the z/OS commands discussed in this appendix.

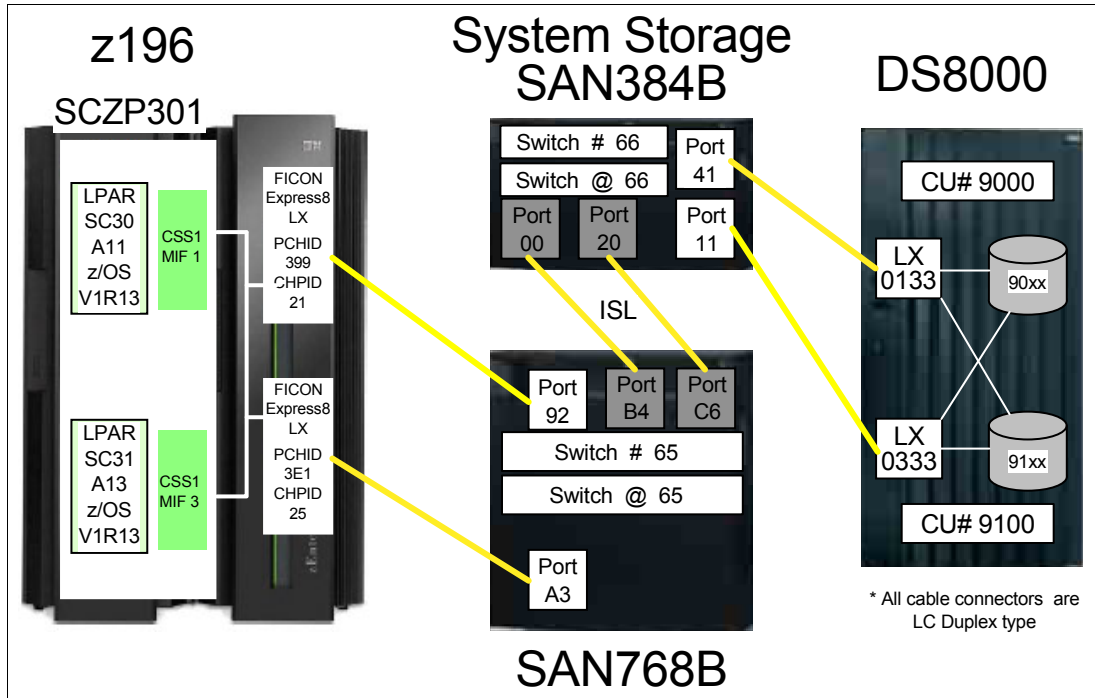


Figure F-1 Our FICON configuration

## Displaying system status using D M=CPU

Example F-1 shows the output of a z/OS D M=CPU command. The output shows the CPU type and serial number. It also defines the CPU characteristics associated to a specific partition (LPAR A11, in this case).

The output shows the General CPs, zIIPs, and zAAPs configured, and the CPU Node Descriptor (ND), the System Information taken from the Store System Information (STSI) instruction, the CPC name, the CSS ID where this partition is running, and its MIF and LPAR IDs.

Example: F-1 Displaying the CPU characteristics

```

D M=CPU
IEE174I 13.25.08 DISPLAY M 007
PROCESSOR STATUS
ID CPU SERIAL
00 + 113BD52817
01 -
02 -A
03 -I
04 -
05 -
06 -A
07 -I
    
```



```

CPC ND = 002817.M32.IBM.02.0000000B3BD5
CPC SI = 2817.716.IBM.02.0000000000B3BD5
      Model: M32
CPC ID = 00
CPC NAME = SCZP301
LP NAME = A11      LP ID = 11
CSS ID = 1
MIF ID = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE

```

```

A      APPLICATION ASSIST PROCESSOR (zAAP)
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

```

---

The other fields in the display output have the following meanings:

► **CPC name**

Reflects the CPC object name customized to the Hardware Management Console (HMC) at System z installation time. The CPC name is also defined in the HCD/HCM (a CPC can only be defined once). HCD/HCM also requires a Proc.ID to be specified for the CPC.

► **LP name**

Is defined in the HCD/HCM logical partition definition or in the IOCP RESOURCE statement. The name must be unique for each logical partition across all four Channel Subsystems (CSSs) on System z. Changing, adding, or deleting LPARs is fully concurrent since System z10.

**Note:** System z10 (2097 and 2098) implemented a fixed Hardware System Area (HSA) storage approach that allows for concurrent adding logical partitions, channel subsystems, subchannel sets, logical processors and cryptographic co-processors. Adding logical processors to a LPAR concurrently requires z/OS 1.10 or later.

► **CSS ID**

Indicates in which CSS this LPAR was created. The number of CSSs that a System z supports depends on the machine. The z196 and z10 Enterprise Class (EC) model support four CSSs in total. The z114 and z10 Business Class (BC) model support two CSS. Increasing or decreasing the number of CSSs in use is concurrent since z10. Each CSS supports three subchannel sets: SS0, SS1, and SS2, except z114 and z10 that only supports two subchannel sets: SS0, and SS1.

► **MIF ID**

A value, between 1 and 15, that is associated to a partition in one CSS. The MIF ID numbers are defined in the HCD/HCM in the IOCP RESOURCE statement.

The MIF ID must be unique for each logical partition in a System z CSS. The MIF ID is used by the System z channel subsystem and channels to identify the source of an I/O request. You need to know the MIF ID to resolve Logical Partition related problems or to resolve failures in establishing a FICON Logical Path.

► LP ID

The Logical Partition ID (LP ID) is specified in the Logical Partition Image Profile using the HMC. The LP ID number, from x'00' to x'3F, must be unique across all partitions on a System z.

## Displaying additional z/OS information using D IPLINFO

Example F-2 displays the output of a z/OS D IPLINFO command.

*Example: F-2 Displaying the z/OS IPL information*

---

```
D IPLINFO
IEE254I 14.51.16 IPLINFO DISPLAY 200
SYSTEM IPLED AT 22.03.35 ON 11/02/2011
RELEASE z/OS 01.13.00 LICENSE = z/OS
USED LOAD01 IN SYS0.IPLPARM ON 0C730
ARCHLVL = 2 MTLSHARE = N
IEASYM LIST = XX
IEASYS LIST = (00) (0P)
IODF DEVICE: ORIGINAL(0C730) CURRENT(0C730)
IPL DEVICE: ORIGINAL(09302) CURRENT(09302) VOLUME(Z1DRB1)
```

---

The output displays the following information:

- The z/OS release (RELEASE z/OS 01.13.00) that this partition is running
- The time stamp, that shows when this partition was IPLed (SYSTEM IPLED AT 14.51.16 ON 11/02/2011)
- The LOAD member selected from the PARMLIB (USED LOAD01 IN SYS0.IPLPARM ON C730)
- The system architecture level (ARCHLVL = 2), indicating a 64-bit operating system in this case
- The Original IODF device, where the I/O configuration resided during system IPL
- The Current IODF device, where the I/O configuration now resides (can be the same)
- The Original IPL device, SYSRES device number from which the system was IPLed
- The Current IPL device, SYSRES current device number (can be the same)
- IPL device VOLSER

**Note:** IPLing the system from the wrong device, or selecting the incorrect IODF device, can lead to I/O-related and configuration-related problems.

Because z/OS 1.11, IODF and IPL device information report the original and the current device number.

## Displaying the I/O configuration using D IOS,CONFIG

Example F-3 displays the following information about the I/O configuration in use by the System z:

- ▶ The IODF data set name (ACTIVE IODF DATA SET = SYS6.IODF09)
- ▶ The active CSS and subchannel sets in use (ACTIVE CSS:1 SUBCHANNEL SETS CONFIGURED:0, 1, 2)
- ▶ The token that will allow dynamic changes to the configuration

A match between the TOKEN IODF and the ACTIVE IODF DATA SET names will indicate whether the system allows dynamic hardware and software changes to the I/O configuration.

*Example: F-3 D IOS,CONFIG z/OS command output*

---

### D IOS,CONFIG

```
IOS506I 15.08.00 I/O CONFIG DATA 246
ACTIVE IODF DATA SET = SYS6.IODF09
CONFIGURATION ID = TEST2094      EDT ID = 01
TOKEN:  PROCESSOR DATE    TIME    DESCRIPTION
SOURCE: SCZP301 11-11-23 11:06:06 SYS6    IODF09
ACTIVE CSS: 1  SUBCHANNEL SETS CONFIGURED: 0, 1, 2
CHANNEL MEASUREMENT BLOCK FACILITY IS ACTIVE
```

---

## Displaying HSA usage using D IOS,CONFIG(HSA)

Example F-4 displays information about the amount of physical control units and subchannels in use per subchannel set in each defined channel subsystem (CSS).

**Note:** The maximum number of subchannels on subchannel set 0 (SS0) is 65280. The maximum number of subchannels on subchannel set 1 and 2 (SS1 and SS2) is 65535.

A *subchannel* corresponds to a device Unit Control Word (UCW), which contains hardware controls and definition information about a particular device.

Note the following points:

- ▶ System z196 and z10 EC can have up to four Channel Subsystems (0,1, 2 and 3) with three subchannel sets on System z196 and two subchannel sets on System z10 EC.
- ▶ System z114 and z10 BC can have up to two Channel Subsystems (0 and 1) with two subchannel sets each.

*Example: F-4 D IOS,CONFIG(HSA) z/OS command output*

---

### D IOS,CONFIG(HSA)

```
IOS506I 15.14.06 I/O CONFIG DATA 260
HARDWARE SYSTEM AREA AVAILABLE FOR CONFIGURATION CHANGES
PHYSICAL CONTROL UNITS          8006
CSS 0 - LOGICAL CONTROL UNITS    3974
SS 0  SUBCHANNELS                47374
SS 1  SUBCHANNELS                61951
SS 2  SUBCHANNELS                65535
CSS 1 - LOGICAL CONTROL UNITS    3954
SS 0  SUBCHANNELS                47092
```

|             |                       |              |
|-------------|-----------------------|--------------|
| <b>SS 1</b> | <b>SUBCHANNELS</b>    | <b>61951</b> |
| <b>SS 2</b> | <b>SUBCHANNELS</b>    | <b>65535</b> |
| CSS 2 -     | LOGICAL CONTROL UNITS | 3952         |
| SS 0        | SUBCHANNELS           | 49115        |
| SS 1        | SUBCHANNELS           | 61951        |
| SS 2        | SUBCHANNELS           | 65535        |
| CSS 3 -     | LOGICAL CONTROL UNITS | 4026         |
| SS 0        | SUBCHANNELS           | 55894        |
| SS 1        | SUBCHANNELS           | 61951        |
| SS 2        | SUBCHANNELS           | 65535        |

---

The output displays:

- ▶ The total number of physical control units defined in the configuration: 8006
- ▶ The number of logical control units per Channel Subsystem
- ▶ The number of subchannels defined in each CSS's subchannel set (0, 1, and 2)

## Display units command D U

Use the z/OS Display Unit command as the starting point for determining the status of a device defined to z/OS, and the z/OS status of the defined paths to the device.

There are a few exceptions to this, where the z/OS Display Unit command will not display the status of a defined device. (DASD ALIAS devices is one of these exceptions; for more information, refer to "Using a D M=DEV(dddd) command for an ALIAS device number" on page 484 and to "DEVSERV command - DS P,dddd,n" on page 484.)

The Display Unit z/OS command can be used to:

- ▶ Determine if the device is defined to z/OS (exception devices are not displayed). If the z/OS status for the requested device number is displayed, it is defined to z/OS.
- ▶ Determine the defined device type: 3390, 3590, CTC, and so on.
- ▶ Identify whether the defined z/OS type is not the same as the device being addressed (when trying to bring the device online); if it is not the same, then the device will not be brought online to z/OS.
- ▶ Indicate the z/OS status of the device. The statuses include:
  - O** Online.
  - F-NRD** Offline.
  - Allocated A** Online and in use by a job or task.
  - A-BSY** Allocated and busy, meaning that there is an I/O operation in progress for the z/OS device in the channel subsystem image (in the channel subsystem for the logical partition that z/OS is running in).
  - VOLSER = NE800** For devices that support "volumes", it indicates which volume is currently mounted on that device.

Whenever a z/OS Display Unit command (or any other general form of this command such as **D U,DASD,,dddd,n** or **D U,,ALLOC,dddd,n**) is entered, z/OS will display the status of the devices that the z/OS System Control Program (SCP) has recorded at that time.

If the devices are not defined to z/OS, or if they are exception devices, z/OS will display one of the following pieces of information:

- ▶ The next device number or the next device number of that type, for a display units type request
- ▶ The next device of that status, for a display units by status
- ▶ If a combination of both type and status is used, then it will be the next device of that type and status (refer to Example F-6 on page 478)

Because of this, always verify that the device (or devices) displayed in response to a z/OS **Display Units** command request is for your display unit's requested device number, or that it includes your requested device number.

## D U,,,dddd,1

The **Display Units** command is useful for determining the z/OS status of a device. The output from this command is shown in Example F-5.

The example shows the normal result of displaying the z/OS status of a group of devices. Devices 9000, 9001, 9002, and 9003 are online and not allocated; that is, the devices are not in use by a job or application.

If the device is online and allocated, you can use the z/OS display allocation command to determine who the device is allocated to. Use the display allocation command, and it will show who the device is allocated to: **D U,,ALLOC,9003,1**

*Example: F-5 z/OS display device status of a group of four devices defined and online*

---

### D U,,,9000,4

```
IEE457I 12.19.16 UNIT STATUS 966
UNIT TYPE STATUS      VOLSER  VOLSTATE
9000 3390 0           LX9000  PRIV/RSDNT
9001 3390 0           LX9001  PRIV/RSDNT
9002 3390 0           LX9002  PRIV/RSDNT
9003 3390 0           LX9003  PRIV/RSDNT
```

---

There are also a few other device statuses that can be returned for a D U command:

|            |   |
|------------|---|
| <b>AS</b>  | Automatically switchable.   |
| <b>BOX</b> | Hardware failure. The device has been BOXED (refer to Boxed status notes:).   |
| <b>BSY</b> | Busy.   |
| <b>C</b>   | Console.  |
| <b>F</b>   | Offline. F indicates that more than one bit is set in the UCB. This is used when a combination of offline and some other status value needs to be displayed (for example, F-NRD). |
| <b>L</b>   | The release on a device is pending and reserve might or might not have occurred.  |
| <b>M</b>   | A device managed by a device manager, such as JES3 or a non-IBM tape management subsystem.  |
| <b>MTP</b> | Mount Pending.  |
| <b>NRD</b> | Not Ready.  |

|                |   |
|----------------|---|
| <b>O</b>       | Online.   |
| <b>OFFLINE</b> | This is used when the only status value that needs to be displayed is OFFLINE.  |
| <b>P</b>       | Reserve Pending.  |
| <b>PND</b>     | Offline pending.  |
| <b>PO</b>      | Offline pending, and also not ready. This status value is displayed only if NRD is also displayed on the status line. |
| <b>PUL</b>     | Unload pending.   |
| <b>R</b>       | Reserved, shared DASD or exclusively-assigned device.   |
| <b>RAL</b>     | Restricted to Allocation.   |
| <b>S</b>       | SYSRES  |
| <b>SPD</b>     | Suspended. (a paging volume). The channel program is temporarily suspended while the system is using the device.      |
| <b>SYS</b>     | Allocated to system.  |
| <b>UNAVL</b>   | The device has been marked as unavailable for allocation by the VARY xxxx,UNAVAIL operator command.                   |

**Boxed status notes:**

- ▶ If the reported status of a device is O-BOX (Online and Boxed), that status can be cleared by using a vary online with unconditional parameter:  
V dddd,ONLINE,UNCOND.
- ▶ If the reported status is F-BOX (Offline and Boxed), the device can be brought back online with the VARY dddd,ONLINE command. This will enable the UCW and perform the online processing to the device. Assuming that the error condition was resolved, the device will come online. If the error condition still exists, however, the device might remain in the boxed state.

**Displaying a device not in current configuration**

If the z/OS display units request is for a device that is not defined in the current OSCONFIG member (in use by z/OS), or if the device is an exception device, then z/OS cannot display the status for that device.

z/OS will display the next displayable device number; or the next device number of the requested device by type (when using the display units type command); or the requested device by status (when using the display unit status command). For this reason, always check that the devices displayed include your requested device, as shown in Example F-6.

*Example: F-6 D U request of a non-existing device*

---

```

D U,,,8F80,1
IEE457I 15.41.25 UNIT STATUS 348
UNIT TYPE STATUS          VOLSER    VOLSTATE
9000 3390 0              LX9000    PRIV/RSDNT

```

---

The D U request shown in this example was for device 8F80; however, notice that the actual display reports the status for device 9000.

## D U,,,ALLOC,dddd,n

The D U,,,ALLOC,dddd,n command is one of the various formats for the Display Unit z/OS device commands. Its main purpose is to display the allocations of a device when its status is determined to be Allocated. Example F-7 shows the device number 9000 as online, allocated to two jobs, and currently in use.

*Example: F-7 z/OS display device allocation*

---

### D U,,,ALLOC,9000,1

```
IEE106I 15.44.47 UNITS ALLOCATED 362
UNIT     JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID   JOBNAME  ASID
9000     DUMPSRV  0005   MVSRECVA 0028
```

---

The information displayed shows that the jobs using the device are:

- ▶ Jobname: DUMPSRV
- ▶ Jobname: MVSRECVA

## V dddd,offline

In Example F-8, the V dddd,OFFLINE command is issued against device 9000. However, as shown in Example F-7, device 9000 is allocated and in use. Therefore, the action of varying this device offline will be in a pending status.

*Example: F-8 Varying an allocated device offline*

---

### V 9000,OFFLINE

```
IEF524I 9000, VOLUME LX9000 PENDING OFFLINE
```

....

### D U,,,9000,4

```
IEE457I 15.49.23 UNIT STATUS 375
UNIT TYPE STATUS VOLSER VOLSTATE
9000 3390 A-PND  LX9000 PRIV/RSDNT
9001 3390 0      LX9001 PRIV/RSDNT
9002 3390 0      LX9002 PRIV/RSDNT
9003 3390 0      LX9003 PRIV/RSDNT
```

---

## Device in OFFLINE status

Example F-9 shows a display reporting that device 9000 is in the OFFLINE state.

*Example: F-9 Device is now offline after the allocation ended*

---

### D U,,,9000,1

```
IEE457I 15.50.38 UNIT STATUS 379
UNIT TYPE STATUS      VOLSER      VOLSTATE
9000 3390 OFFLINE      /RSDNT
```

---

A device can be offline for a number of reasons:

- ▶ It has been varied offline by the operator (case shown in Example F-9).
- ▶ It was defined to the operating system as being offline at system IPL time.
- ▶ The device is defined to the operating system but not to the logical partition that the z/OS is running in; that is, there is no subchannel.
- ▶ The device is defined to the operating system but not to the channel subsystem; that is, there is no subchannel.

- ▶ There is no operational path to the control-unit or device.
- ▶ All the operational paths were logically varied offline.
- ▶ The definition to access the control unit device is not correct.
- ▶ CHPID.
- ▶ Link destination link address (destination port address).
- ▶ CUADD - the CU logical image address.
- ▶ For an IBM 2107, the CUADD is unique for an LSS.
- ▶ For a CTC, the CUADD is the MIF ID of the target LPAR.
- ▶ The device addressing definition is not correct.
- ▶ The device is not ready.
- ▶ Link destination link address (destination port address).
- ▶ The device is in an error state.
- ▶ The defined device type and the physical device type (the actual device) do not match.
- ▶ An error condition occurred during the use of the device (that is, a reserve or assign was lost) and the device was forced offline (it might have been boxed during the recovery process, as well).

### V dddd,ONLINE

Use the z/OS command shown in Example F-10 to bring a single device online. There are variations to the V dddd,ONLINE command that allow more than one device to be brought online simultaneously. A group of devices can be varied online using dddd,dddd or dddd-dddd as variables for the vary command.

*Example: F-10 V dddd, ONLINE results*

---

```
V 9000,ONLINE
IEE302I 9000      ONLINE
```

---

A failure to bring a device online can be caused by:

- ▶ There is no UCB for the target device.
- ▶ There is no UCW (subchannel) for the target device.
- ▶ All paths to the device are offline.
- ▶ None of the defined paths are operational.
- ▶ The CU and or the device addressing is not correct.
- ▶ There is a mismatch of device type between the definition and the actual device.

### D M=DEV(dddd)

The D M=DEV(dddd) command displays the device status for device dddd. Example F-11 shows the output of this display device matrix command. It shows the path status of all paths that have been defined to this single device address (device 9000, in the example). The destination link address and the entry link address are individually provided for each channel accessing the device.

*Example: F-11 z/OS display output for D M=DEV(dddd)*

---

```
D M=DEV(9000)
IEE174I 15.56.20 DISPLAY M 396
DEVICE 9000  STATUS=ONLINE
```



```

CHP                40  41  42  43  21  25
ENTRY LINK ADDRESS 22  22  1E  1E  6592 65A3
DEST LINK ADDRESS  1B  1B  10  10  6641 6611
PATH ONLINE        Y   Y   Y   Y   Y   Y
CHP PHYSICALLY ONLINE Y   Y   Y   Y   Y   Y
PATH OPERATIONAL   Y   Y   Y   Y   Y   Y
MANAGED            N   N   N   N   N   N
CU NUMBER          9000 9000 9000 9000 9000 9000
MAXIMUM MANAGED CHPID(S) ALLOWED: 0
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.932.IBM.75.0000000L3001.0000
SCP TOKEN NED      = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED     = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW, ZHPF

```

---

Example F-12 is a variation of the D M=DEV(dddd) command where a unique CHPID that accesses the device dddd is selected for display.

*Example: F-12 Variation of D M=DEV(dddd,(cc))*

---

```

D M=DEV(9000, (21))
IEE174I 16.01.05 DISPLAY M 407
DEVICE 9000 STATUS=ONLINE
CHP                21
ENTRY LINK ADDRESS 6592
DEST LINK ADDRESS 6641
PATH ONLINE        Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL   Y
MANAGED            N
CU NUMBER          9000
DESTINATION CU LOGICAL ADDRESS = 00
SCP CU ND          = 002107.932.IBM.75.0000000L3001.0000
ACTUAL CU ND       = 002107.900.IBM.75.0000000L3001.0133
SCP TOKEN NED      = 002107.900.IBM.75.0000000L3001.0000
ACTUAL TOKEN NED   = 002107.900.IBM.75.0000000L3001.0000
SCP DEVICE NED     = 002107.900.IBM.75.0000000L3001.0000
ACTUAL DEVICE NED = 002107.900.IBM.75.0000000L3001.0000
HYPERPAV ALIASES CONFIGURED = 106
FUNCTIONS ENABLED = MIDAW, ZHPF

```

---

For both D M=DEV examples, the following important information is displayed:

- ▶ The device current status: DEVICE 9000 STATUS=ONLINE.
- ▶ The FICON Director entry port: 6592 - FICON Director @ 65, Port 92.
- ▶ The FICON Director exit port: 6641 - FICON Director @ 66, Port 41.
- ▶ Note that the FICON Director addresses are different for the entry and exit (Dest) Ports. This is an indication that a FICON Director cascading topology is being used.
- ▶ PATH ONLINE = Y.
- ▶ The PATH ONLINE field indicates whether the path is logically online to z/OS and will change as a result of the **VARY PATH** command. Information for this field is obtained from the UCB Logical Path Mask (UCBLPM).

- ▶ CHP PHYSICALLY ONLINE = Y.
- ▶ The CHP PHYSICALLY ONLINE field indicates whether the path is physically available. This field will change as a result of the z/OS CF CHP command. Information for this field is obtained from the Path Available Mask (PAM) in the Unit Control Word (UCW) or subchannel.

**Note:** The CHP PHYSICALLY ONLINE field might not be accurate when the CHPID has been “stolen” from z/OS using the HMC/SE CHPID Config OFF/ON facility instead of the z/OS CF CHP(cc),OFFLINE command.

- ▶ PATH OPERATIONAL = Y

If PATH OPERATIONAL=N means that the PATH OPERATIONAL field status is changed as a result of the channel subsystem attempting an I/O operation on a path and the path responding with a not operational I/O interface sequence. This response can be due to:

- Disabling a channel interface Host Bay Adapter (HBA) at the control unit
- Disabling a port on a FICON Director
- Powering off a control unit or cluster
- Powering off a DASD controller

**Note:** The PATH OPERATIONAL status can be easily misinterpreted. The information displayed is a direct reflection of a path mask called the Path Operational Mask (POM) in the UCW.

- ▶ PATH NOT VALIDATED = Y

The PATH NOT VALIDATED field might be shown if the device was not online at IPL time. If a device is not online when the IPL device pathing process occurs, the UCBVALPH bit is set to On to indicate that the paths to the device were not validated. Validation will occur when the device is varied online.

- ▶ MANAGED = N

This field indicates whether the channel is using the Dynamic Channel Management (DCM) facility of z/OS.

- ▶ CU NUMBER = 9000

This field presents the Control Unit Number associated to the physically attached defined CU.

- ▶ DESTINATION CU LOGICAL ADDRESS= 00

The DESTINATION CU LOGICAL ADDRESS field shows the logical link address of a control unit that is used to access the I/O device associated with the specified channel paths.

- ▶ SCP CU ND = 002107.932.IBM.75.0000000L3001.0000

This shows the Node Descriptor (ND) LAST obtained by the System Control Program (SCP).

- ▶ ACTUAL CU ND = 002107.900.IBM.75.0000000L3001.0133

This shows the Node Descriptor of the attached subsystem read from the device in response to this command.

- ▶ SCP TOKEN NED = 002107.900.IBM.75.0000000L3001.0000

This shows the Node-Element Descriptor last obtained by the SCP.

- ▶ ACTUAL TOKEN NED = 002107.900.IBM.75.0000000L3001.0000  
This shows the Node-Element Descriptor read from the device in response to this command.
- ▶ SCP DEVICE NED = 002107.900.IBM.75.0000000L3001.0000  
This shows the Node-Element Descriptor last obtained by the SCP.
- ▶ ACTUAL DEVICE NED = 002107.900.IBM.75.0000000L3001.0000  
This shows the Node-Element Descriptor read from the device in response to this command.
- ▶ RNID = xxxxx.xxx.xxx.xx.xxxxxxxxxx.xxxx (not shown in the example)  
This is the Control Unit Remote Node IDentifier obtained by the Channel Subsystem during Channel - CU initialization. This RNID information is kept in the Hardware System Area (HSA).
- ▶ HYPERPAV ALIASES CONFIGURED = 106  
This field shows the number of (ALIAS) HYPERPAV ALIAS devices configured in this LCU. See Example F-14 on page 484 for a display of an alias device through the D M=DEV command.
- ▶ FUNCTIONS ENABLED = MIDAW, ZHPF  
This field shows the additional facilities that this device is capable of using.

### Using a D M=DEV command without specifying a device

Example F-13 shows the output of a display device (D M=DEV) command without specifying a device number. The information for each defined device will be provided by the display, including device-specific information about its current state.

*Example: F-13 Using D M=DEV without specifying a device*

---

```

D M=DEV
IEE174I 16.16.53 DISPLAY M 453
DEVICE STATUS: NUMBER OF ONLINE CHANNEL PATHS
      0 1 2 3 4 5 6 7 8 9 A B C D E F
02EA 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
0413 DN DN DN DN . . . . DN DN DN DN 1 1 1 1
04B0 1 1 1 1 . . . . 1 1 1 1 . . . .
0513 DN DN DN DN 1 1 1 1 DN DN DN DN # # # #
05B0 1 1 1 1 . . . . 1 1 1 1 . . . .
0600 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4
0601 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4 4
....
19E9 UL UL UL UL UL UL UL UL UL UL UL UL UL UL HA HA
19EA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
....
19EF HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA HA
***** SYMBOL EXPLANATIONS *****
@ ONLINE, PHYSICALLY ONLINE, AND OPERATIONAL INDICATORS ARE NOT EQUAL
+ ONLINE          # DEVICE OFFLINE          . DOES NOT EXIST
BX DEVICE IS BOXED          SN SUBCHANNEL NOT AVAILABLE
DN DEVICE NOT AVAILABLE     PE SUBCHANNEL IN PERMANENT ERROR
AL DEVICE IS AN ALIAS       UL DEVICE IS AN UNBOUND ALIAS
HA DEVICE IS A HYPERPAV ALIAS HU HYPERPAV ALIAS UNUSABLE

```

---

'D M=DEV', which is shown in Example F-13 on page 483, provides symbols identifying the state of existing devices. Significant symbolic meanings are explained here:

- DN** The DN symbol means that the device has been defined to the z/OS IOGEN (OSCONFIG) but has *not* been defined in the IOCP; that is, there is no subchannel for this device known to the channel subsystem. This normally occurs when the wrong IOCDS or OSMCONFIG is loaded. The UCB for a DN device has the Not Connected bit left on during the device mapping time during IPL.
- .** The dot (.) indicates that the device does not exist or has not been defined.
- 1** The numeric digit (1) indicates the number of paths defined for the device. It is possible to have a combination of a digit and a symbol. For example, 4@ indicates an out-of-line path condition; online, physically online, and operational indicators are not equal with at least one of these paths, where:
- Online means the path is online.
  - Physically online means the CHPID is physically online.
  - Operational means the path is operational (POM).

### Using a D M=DEV(dddd) command for an ALIAS device number

Example F-14 shows an output of a D M=DEV(dddd) issued against an ALIAS device address.

*Example: F-14 Displaying an ALIAS device using D M=DEV(dddd)*

---

```
D M=DEV(9080)
IEE174I 18.22.19 DISPLAY M 660
DEVICE 9080 STATUS=HYPERPAV ALIAS
```

---

### Using a D M=DEV(dddd) command to a non-existing device

Example F-15 shows the output of D M=DEV(dddd) issued against a non-existing device. There is no UCB available for this device in the OSMCONFIG being used.

*Example: F-15 Displaying a non-existing device using 'D M=DEV(dddd)*

---

```
D M=DEV(4000)
IEE174I 18.23.52 DISPLAY M 666
DEVICE 4000 STATUS=NOT IN SYSTEM
```

---

## DEVSERV command - DS P,dddd,n

The DEVSERV command requests a display of the current status of DASD and TAPE devices.

The response is a display of basic status information about a device, a group of devices, or storage control units, and optionally can include a broad range of additional information. With the DEVSERV command you can display:

- ▶ Device Number
- ▶ Device type
- ▶ Logical mode of the device

- ▶ Number of data sets allocated on the volume
- ▶ Volume serial label
- ▶ Channel path ID
- ▶ Status of the path
- ▶ Status of an SMS-managed device:
  - Volume Status
  - Storage Group name
  - Storage Group status
- ▶ Control unit type and model
- ▶ Control unit serial number
- ▶ Device capacity, in cylinders
- ▶ Device extended function status
- ▶ Unit Control Block (UCB) device type information
- ▶ Help text, when you request it
- ▶ The following information can be displayed, if the device belongs to the DASD storage subsystem:
  - Real device type (if what is shown is an emulated device type)
  - Control unit type and model (or emulated control unit type and model, if the real and emulated control units are not the same)
  - Subsystem ID for this storage subsystem
  - Cache fast write state
  - Track caching state
  - DASD fast write state
  - State of pinned data
  - State of dual copy, PPRC, or SPARing (if there is any)
  - Address of the other device in a dual copy pair
  - Channel subsystem device address
  - Subsystem internal logical device address
  - An indication if the device extended function status information is inconsistent between z/OS control blocks and the storage subsystem
  - An indication if the defined (UCB) device type is inconsistent with the real device type
  - Optionally, the total number of cylinders for each unique track format (3380, 3390, and 9345) for all of the devices within the scope of the request
- ▶ The following information occurs if the device belongs to a tape library:
  - Device type equivalent to DTYPE from the DS P command
  - Device status indicating online/offline and ready/not ready
  - Device type and model
  - Device serial number
  - Library identification number
  - An indication if the defined (UCB) device type is inconsistent with the real device type

Example F-16 on page 486 shows the output of a DEVSERV command. The response is a display of basic status information about a device, a group of devices, or storage control units,

and optionally can include a broad range of additional information. The path status for all defined channel paths is included, as well.

There are different ways of using the DEVSERV command:

- ▶ DS P,1C00: Test device 1C00 only through all paths.
- ▶ DS P,2000,16,ON: Test only ONLINE devices in 2000-200F.
- ▶ DS QD,3200,1,UCB,DCE,SSSCB,DPCT: Display MVS control blocks (1 device).
- ▶ DS QD,4800,1,RDC,RCD,SNSS: Get hardware information and update control blocks.

*Example: F-16 Devserv QDASD with options RDC and DCE*

---

**DS QD,9000,RDC,DCE**

```
IEE459I 16.38.57 DEVSERV QDASD 516
UNIT VOLSER SCUTYPE DEVTYPE      CYL  SSID SCU-SERIAL DEV-SERIAL EFC
09000 LX9000 2107932 2107900      10017  8940 0175-L3001 0175-L3001 *OK
  READ DEVICE CHARACTERISTIC
2107E833900C5E8C 7FF720322721000F E000E5A205940222 1309067400000000
0000000000000000 32321F02DFEE0001 0677080F007F4A00 003C000000002721
  DCE AT V00EA01F8
487880712B371CF0 0000000005E674E0 D800272127203232 1FF70823016D0000
00FD59F49400F07E 001F3C1E00038C00 0000000000002721 0000272000000000
0000000000000000
****          1 DEVICE(S) MET THE SELECTION CRITERIA
****          0 DEVICE(S) FAILED EXTENDED FUNCTION CHECKING
```

---

- ▶ DS QT,TYPE=3490<,DEFINED>: Display TAPE with 3490 information. You can use the QTAPE (QT) parameter to display tape information. Using the MED (medium) option DS QT,xxxx,MED,nnn allows you to display information for the device type, media type, and cartridge volume serial number.
- ▶ DS QPAVS,dddd,n: Display status of Parallel Access Volume devices.

To learn about other uses of the DEVSERV command, refer to *z/OS MVS System Commands, SA22-7627*, or to *DFSMS Storage Administration Reference, SC26-7402*.

**Note:** The DEVSERV command is one of the few commands that actually performs an I/O operation against the device or devices specified. Most display commands do not actually perform a real I/O operation and report back the state of specific fields in the UCB and UCW (subchannel).

Example F-17 shows the DS P,dddd,nnn command.

*Example: F-17 Output of one type of DEVSERV command: DS P,dddd,nnn*

---

**DS P,9000,4**

```
IEE459I 12.00.02 DEVSERV PATHS 949
UNIT DTYPE  M CNT VOLSER  CHPID=PATH STATUS
      RTYPE  SSID CFW TC   DFW PIN DC-STATE CCA DDC      CYL CU-TYPE
09000,33909 ,0,000,LX9000,40=* 41=+ 42=+ 43=+ 21=+ 25=+ 22=+ 26=+
      2107   8940 Y YY. YY. N SIMPLEX 00 00      10017 2107
09001,33909 ,0,000,LX9001,40=+ 41=+ 42=+ 43=+ 21=+ 25=+ 22=+ 26=+
      2107   8940 Y YY. YY. N SIMPLEX 01 01      10017 2107
09002,33909 ,0,000,LX9002,40=+ 41=+ 42=+ 43=+ 21=+ 25=+ 22=+ 26=+
      2107   8940 Y YY. YY. N SIMPLEX 02 02      10017 2107
09003,33909 ,0,000,LX9003,40=+ 41=+ 42=+ 43=+ 21=+ 25=+ 22=+ 26=+
      2107   8940 Y YY. YY. N SIMPLEX 03 03      10017 2107
```

```
***** SYMBOL DEFINITIONS *****  
O = ONLINE                + = PATH AVAILABLE  
* = LOGICALLY OFF, PHYSICALLY ON
```

---

The fields reported by the DASD DEVSERV command are listed here:

▶ DTYPE = 33909

This is the device type of the device reported by DEVSERV command. In this case, device 9000 reported back as a 3390 model 9.

▶ RTYPE = 2107

This is the real (true) device type of the device reported by the DEVSERV command. 2107 is the true device type where the DEVSERV command has been executed.

▶ M = O

This field represents the Device UCB status:

A - Allocated  
F - Offline  
M - Mount Pending  
O - Online  
P - Offline Pending  
N - Not Allocated

▶ CNT

This field indicates the number of data sets allocated on this volume.

▶ SSID

This field indicates the subsystem ID of the device CU.

▶ VOLSER

This field indicates the VOLSER of the device pointed to by the **DEVSERV** command.

▶ CFW

Cache Fast Write, which indicates the status of the CFW:

Y - CFW is Active  
N - CFW is Inactive  
S - CFW is Suspended with Pinned Data

▶ TC (c,d) C=Device Cache; D=Subsystem Cache

These two characters indicate the Device Cache and Subsystem Cache status:

Y - Active  
N - Inactive  
A - Pending Active  
F - Pending Inactive  
M - Disabled for maintenance (this will override other statuses)  
P - Pending Inactive - Destage in Progress  
S - CU is running in Single Cluster mode (half of the cache in one cluster is not in use)  
T - Terminated due to a Subsystem Cache Error

▶ DFW (e,f) E=DASD Fast Write; F=NVS Status

These two characters indicate the DASD Fast Write and the NVS status:

Y - Active/Available  
N - Inactive/Unavailable  
F - NVS Pending Unavailable/Disable - Destage failed

I - DFW - Deactivate Pending - Destage in Progress  
 I - NVS - Battery Defective  
 P - NVS Pending Unavailable - Destage in Progress or has failed  
 S - DFW is temporarily Suspended with Pinned Data  
 U - DFW Deactivate Pending with Destage failed  
 U - NVS Terminated due to error

▶ PIN (p)

This single character indicates whether pinned data exists:

N - No Pinned Data  
 Y - Pinned Data in Cache or NVS - CFW/DFW allowed  
 S - Retriable Pinned Data in Cache/NVS - CFW/DFW temporarily suspended

▶ DC-STATE

This field reports the Dual Copy current status:

SIMPLEX (Simplex: Not Duplex-Pair)  
 PRIMARY (Primary Device of Active Duplex-Pair)  
 SECONDARY (Secondary Device of Active Duplex-Pair)  
 PRI-PNDG (Establish/Copy in Progress as Primary)  
 SEC-PNDG (Establish/Copy in Progress as Secondary)  
 PRI-SDPL (Suspended, current and original Primary) - DDC same)  
 SEC-SDPL (Suspended, current and original Secondary) - DDC same)  
 PRI-SSEC (Suspended, was Secondary and now Primary - DDC Swapped)  
 SEC-SPRI (Suspended, was Primary and now Secondary - DDC Swapped)  
 SPARE (RAMAC: RDC byte 57=33/34 and SNSS byte 36 bit 4,5=01/10)  
 SPAR-PNDG (Spare Device being copied or has been copied and Copy-Back is pending)  
 SPAR-BRKN (Broken Device copied to Spare device)  
 PPRIMARY (PPRC Primary) PPRI-PNDG (PPRC Primary Pending)  
 PPRI-FAIL (PPRC Primary Fail) PPRI-SUSP (PPRC Primary Suspended)  
 PSECONDARY (PPRC Secondary) PSEC-PNDG (PPRC Secondary Pending)  
 PSEC-FAIL (PPRC Secondary Fail) PSEC-SUSP (PPRC Secondary Suspend)  
 MIRR-OPERIPENDIFAIL (mirroring status per SNSS byte 26 bit 6, 7)

▶ CCA-xx

xx - This indicates the Channel Connection address:  
 – The xx value should match UNITADD-XX in the IOCP.

▶ DDC-yy

yy - This indicates the Device-to-Director Connection address.

▶ CYL

This field reports the number of cylinders of the subject device, when applicable.

▶ CU-TYPE

This field reports the control unit type of the subject device.





## Adding FICON CTC connections

We offer the following recommendations for adding CTC connections to a FICON environment.

You can choose one of the following FICON channel CTC configurations:

- ▶ One FICON channel in a point-to-point, switched, or cascaded configuration
- ▶ Two FICON channels in a point-to-point, switched, or cascaded configuration

The channel operates in full duplex mode, sending and receiving data concurrently.

The FICON channel CTC communication does not require a pair of channels because it can communicate with any FICON channel that has a corresponding FCTC control unit defined.

Although a single FICON native channel per server can provide CTC connections across multiple LPARs, for large FICON configurations we recommend using at least one pair of FICON native channels. Using a pair of FICON native channels allows the installation to maintain the same CTC device definition methodology for FICON as was previously used for ESCON. But the FICON channels can support the definition of FCTC control units and other I/O control units, such as disk and tape, simultaneously.

The FICON channel at each end of the FICON CTC connection, supporting the FCTC control units, can also communicate with other FICON native control units, such as disk and tape connected using a FICON Director.

Because the FICON native channel supports a larger number of devices, installations with a high number of logical partitions in an FCTC complex are easier to design and define.

With the System z, a single FICON CHPID can be used simultaneously to talk to another FICON CHPID and any other native FICON controller.

In a mixed CTC/controller configuration, performance levels can be kept at a premium if the channel utilization is kept at recommended levels (under 50%).

## FCTC device numbering scheme

There are a number of schemes that can be used for developing the device numbers for the FICON CTC devices. Defining a FICON CTC configuration requires an understanding of CTC operations and the different methods available for defining the connections. It is also necessary to take into consideration the subsystem requirements for using CTC connections. As the number of images in a client environment grows, the CTC definition process becomes increasingly complex.

Using the FICON sender/receiver device numbering scheme is the recommended approach. Apart from ease of migration, operations and systems programmers can easily identify the use (send or receive) and target system for any given CTC device. It also simplifies the CTC definition process in that the same operating system CTC device definitions can be used by all z/OS operating systems, independent of the image and processor where that operating system runs.

Using two FICON (FC) channels per server, define all send FCTC control units and devices on one FICON (FC) channel, and define all receive FCTC control units and devices on a second FC channel.

The method makes use of the 4-digit z/Architecture device number where:

► First digit:

- An even hexadecimal number is used for the send CTC control unit and device.
- An odd hexadecimal number is used for the receive CTC control unit and device.

► Second and third digits

These digits represent an assigned CTC image-ID or logical partition identifier (that is, assigned on paper) for the LPAR image. It can be CSS and MIFID or it can be any unique value within the CTC complex. This CTC image-ID/logical partition identifier is used as a target identifier (for the image) that you use when defining how all other images access this image.

► Fourth digit

This digit is set to 0 (zero) for the last number for the control unit and then zero to “F” (0 - F) for the devices and is used to indicate whether the CTC connection is the primary or alternate (backup for availability purposes). A value of 0 to 7 is used for the primary CTC connection devices, and a value of 8 to F is used for the alternate CTC connection devices.

The following example translates into primary devices 4120-4127 and alternate devices 4128-412F on CU 4120:

CU: 4120 represents CPC=4, CTC/LPid=12, CU=0.

Dev: 4120-4127 represents CPC=4, CTC/LPid=12, devices 4120-4127 (primary devices)

Dev: 4128-412F represents CPC=4, CTC/LPid=12, devices 4128-412F (alternate devices)

**Note:** In IOCP on the control unit (CNTLUNIT) statement and the IODEVICE statement you must specify the same unit address (UNITADD=) for devices from different logical partitions (LPARs) to be able to communicate with each other.

For shared FC channels, on the CNTLUNIT statement, the CUADD must be the *destination* control unit multiple image facility ID (MIFID). The MIFID is the logical CSS ID prefixed to the logical partition number (for example; CSS0 LPAR 4 would have CUADD=04 whereas CSS1 LPAR 2 would have CUADD=12). Also, when attached to an unshared FC channel path, the logical address (CUADD) must be zero or not specified.

To fully exploit the full duplex data flow capabilities of the FICON channels, a further performance enhancement can be made. However, this requires a large amount of customization and might not be practical with large configurations. Using two FICON channels per server, evenly spread the send CU/device definitions across the two FC channels, and evenly spread the receive CU/device definitions across the two FC channels.

We do not recommend using the matching device number scheme because it does not enable the operator or systems programmer to identify the use or target system for a given CTC device number. It also makes it difficult to relocate operating systems between logical partitions on the same server or a different server.

For more details, refer to:

- ▶ *FICON CTC Implementation*, REDP-0158



# 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

For information about ordering these publications, see “How to get Redbooks” on page 495. Note that some of the documents referenced here might be available in softcopy only:

- ▶ *IBM System z Connectivity Handbook*, SG24-5444
- ▶ *IBM zEnterprise System Technical Introduction*, SG24-7832
- ▶ *IBM zEnterprise 196 Technical Guide*, SG24-7833
- ▶ *IBM zEnterprise 196 Configuration Setup*, SG24-7834
- ▶ *IBM zEnterprise 114 Technical Guide*, SG24-7954
- ▶ *Getting Started with the IBM 2109 M12 FICON Director*, SG24-6089
- ▶ *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *Implementing the Cisco MDS9000 in an Intermix FCP, FCIP, and FICON Environment*, SG24-6397
- ▶ *FICON Implementation Guide*, SG24-6497
- ▶ *IBM System Storage DS8000 Architecture and Implementation*, SG24-6786
- ▶ *Getting Started with the McDATA Interpid FICON Director*, SG24-6857
- ▶ *Getting Started with the INRANGE FC/9000 FICON Director*, SG24-6858
- ▶ *IBM Tivoli System Automation for z/OS Enterprise Automation*, SG24-7308
- ▶ *IBM System z10 Enterprise Class Technical Introduction*, SG24-7515
- ▶ *IBM System z10 Enterprise Class technical Guide*, SG24-7516
- ▶ *IBM/Cisco Multiprotocol Routing: An Introduction and Implementation*, SG24-7543
- ▶ *IBM System Storage/Brocade Multiprotocol Routing: An Introduction and Implementation*, SG24-7544
- ▶ *Implementing an IBM/Cisco SAN*, SG24-7545
- ▶ *IBM System z10 Enterprise Class Configuration Setup*, SG24-7571
- ▶ *FICON CTC Implementation*, REDP-0158
- ▶ *Disk Storage Access with DB2 for z/OS*, REDP-4187
- ▶ *How Does the MIDAW Facility Improve the Performance of FICON Channels Using DB2 and Other Workloads?*, REDP-4201
- ▶ *Multiple Subchannel Sets: An Implementation View*, REDP-4387
- ▶ *Cisco FICON Basic Implementation*, REDP-4392

## Other publications

These publications are also relevant as further information sources:

- ▶ *z/OS MVS Diagnosis: Reference*, GA22-7588
- ▶ *z/OS MVS Diagnosis: Tools and Service Aids*, GA22-7589
- ▶ *System z Planning for Fiber Optic Links (ESCON, FICON, InfiniBand, Coupling Links, and Open System Adapters)*, GA23-0367
- ▶ *IBM System Storage SAN768B Installation, Service, and User's Guide*, GA32-0574
- ▶ *IBM System Storage SAN384B Installation, Service, and User's Guide*, GA52-1333
- ▶ *IBM System Storage DS8000 Introduction and Planning Guide*, GC35-0515
- ▶ *FICON Express2 Channel Performance Version 1.0*, GM13-0702
- ▶ *System z10 Enterprise Class System Overview*, SA22-1084
- ▶ *z/OS MVS System Commands*, SA22-7627
- ▶ *z/OS MVS System Messages*, SA22-7637
- ▶ *z/Architecture, Principles of Operation*, SA22-7832
- ▶ *System z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037
- ▶ *System z10 Processor resource/Systems Manager Planning Guide*, SB10-7153
- ▶ *DFSMS Storage Administration Reference*, SC26-7402
- ▶ *HMC Operations Guide*, SC28-6830
- ▶ *HMC Operations Guide*, SC28-6873
- ▶ *System z10 Support Element Operations Guide*, SC28-6879
- ▶ *HCD User's Guide*, SC33-7988
- ▶ *z/OS RMF User's Guide*, SC33-7990
- ▶ *z/OS RMF Report Analysis*, SC33-7991
- ▶ *z/OS RMF Performance Management Guide*, SC33-7992
- ▶ *z/OS RMF Programmer's Guide*, SC33-7994
- ▶ *System z Maintenance Information for Fiber Optic Links (ESCON, FICON, Coupling Links, and Open System Adapters)*, SY27-2597
- ▶ *IBM Fiber Optic Cleaning Procedure*, SY27-2604
- ▶ *z/OS RMF Reference Summary*, SX33-9033
- ▶ *Brocade Fabric OS Administrator's Guide*, 53-1001185
- ▶ *High Performance FICON for System z Technical Summary for Customer Planning*, ZSW03058USEN
- ▶ *Performance Considerations for a Cascaded FICON Director Environment Version 0.2x*, Richard Basener and Catherine Cronin  
<http://www-03.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130237.pdf>
- ▶ *IBM System z9 I/O and FICON Express4 Channel Performance*, ZSW03005USEN
- ▶ *IBM System z10 I/O and High Performance FICON for System z Channel Performance*, ZSW03059USEN
- ▶ *IBM System Storage SAN768B*, TSD03037USEN
- ▶ *IBM System Storage SAN768B Fiber Backbone Interoperability Matrix*

- ▶ *Cisco MDS 9506 for IBM System Storage*, TSD00069USEN
- ▶ *Cisco MDS 9509 for IBM System Storage*, TSD00070USEN
- ▶ *Cisco MDS 9513 for IBM System Storage*, TSD01754USEN
- ▶ *Cisco MDS 9506, 9509, 9513 for IBM System Storage Directors Interoperability Matrix*

## Online resources

These Web sites are also relevant as further information sources:

- ▶ Fibre Channel standard Web site  
<http://www.t11.org>
- ▶ Brocade Communications Systems, Inc. Web site  
<http://www.brocade.com>
- ▶ Cisco Systems, Inc. Web site  
<http://www.cisco.com>

## How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)





# Index

## Numerics

2817DEVICE PSP bucket 122, 158, 201  
50µm or 62.5µm multi mode (MM) 159, 201  
50µm or 62.5µm multi-mode (MM) 122  
9µm single mode (SM) 122, 159, 201

## A

Accessible Devices 363  
ACTIVATE IODF=05,TEST 141  
ACTIVATE IODF=06,TEST 183  
ACTIVATE IODF=07,TEST 231  
Activate the production IODF 141, 182, 230  
ACTIVATE,IODF=05,SOFT,TEST 142  
ACTIVATE,IODF=06,SOFT,TEST 183  
ACTIVATE,IODF=07,SOFT,TEST 231  
active CSS 475  
Active Zone  
    Config 295  
ADDRESS keyword 135, 176, 220  
Advanced Performance  
    Tuning Policy 279  
Allow/Prohibit Matrix 355  
American National Standard Institute (ANSI) 4  
Analyze channel information 151, 194, 242  
Analyze Channel Information panel (SE)  
    CSS-to-tag relationship 442  
    Node Descriptor (ND) 441  
Analyze Channel Information panels  
    Tag 441  
Analyze Channel Information panels (SE) 440  
Analyze Control Unit Header PD panel  
    Search by Device Number 444  
    Tag information 445  
Analyze Device Status PD panel  
    Link and Control Unit Address 447  
Analyze FICON problems 370  
Analyze Link Error Statistics Block PD panel  
    PPE 450  
Analyze Paths to a Device PD panel  
    Cuadd 446  
    device path information 446  
    Linkaddr 446  
    subchannel 446  
    Switch Number 446  
Analyze Serial Link Status 153, 195, 243  
Analyze Serial Link Status PD panel  
    Channel link address 448  
    logical pathing information 448  
    Switch number 448  
    Switch number valid 448  
    typical link error messages 448  
Analyze Subchannel Information PD panel  
    Device Number 442  
    Irpt parm 443

LPE 444  
LPM 443  
LPUM 444  
PAM 444  
Pathing information 443  
PIM 444  
PNOM 444  
POM 444  
    subchannel number 442  
    unit address (UA) 442  
ANSI standard architecture 360  
ANSI standard for Virtual Fabrics 101  
APARS 80  
Assign/Unassign Device to Esoteric 127, 165, 207  
attached FICON Director  
    error condition 56  
attached node (0203) 152  
attached node (6500) 194  
attached node (6540) 242  
auditcfg, CLI command 375  
auditdump, CLI command 365, 375

## B

Bandwidth management 108  
BM HyperPAV 101  
BM zCP3000 study 111  
bottleneckmon, CLI command 376  
BOX\_LP 384  
BOX\_LP=(ALL) 384  
buffer credit 34, 51, 104, 246, 328  
buffer-to-buffer 34  
Buffer-to-buffer credits 34, 84  
Build and manage S/390 microprocessor IOCDs and  
IPL attributes 138, 180, 228  
Build Production I/O Definition File 138, 179, 227  
building a production IODF 137, 179, 227  
Business Class (BC) 473

## C

Cascaded 23  
Cascaded (through two FICON Directors) 14  
cascaded FICON Channel-to-Channel (FCTC) 220  
cascaded FICON Director  
    Additional configuration steps 262  
cascaded topology 197–198, 409  
CBDA674I IOCP successfully completed 140  
CBDQAJSK - HCM-HCD Agent Jcl SKELETON 410  
CBDQDISP - HCM-HCD DISPATCHER 410  
CCW 26  
    other parameters 26  
CCW chaining 322  
CCW execution 28  
    control unit 28  
Central Processor Complex (CPC) 320, 409, 472

- CF CHP 147, 482
- CF CHP(20),ON 147
- CF CHP(21),ON 237
- CF CHP(22),ON 189
- Change Channel Path Definition 123, 160
- Change Control Unit Definition 125, 163
- changing passwords on Director 260
- channel 377
- CHANNEL and CONTROL UNIT LOGOUT DATA 360
- Channel Control Check (CCC) 354, 378
- Channel Data Check (CDC) 378
- Channel Detected Errors 352, 354
- Channel feature 37
- Channel Hardware Failure
  - I/O Trace 461
- Channel Information
  - window 194, 242
- channel node (20) 152
- channel node (21) 242
- channel node (22) 194
- channel path (CP) 17, 22, 147, 189, 237, 353, 383, 409
- Channel Path ID . Link Address 125, 164, 206
- Channel Path Identifier 17
- Channel Path List 162, 204
- channel paths, CUs (CPC) 352
- Channel Problem Determination 151
- Channel Problem Determination (CHPD) task panels (SE) 438
  - Analyze Channel Information panels 440
  - Analyze Control Unit Header PD panel 444
  - Analyze Device Status PD panel 446
  - Analyze Link Error Statistics Block (LESB) PD panel 450
  - Analyze Paths to a Device PD panel 445
  - Analyze Serial Link Status PD panel 447
  - Analyze Subchannel Information PD panel 442
  - Optical Power Measurement panel 450
- Channel spanning 18
- Channel Subsystem 17, 98, 476
  - channel path 414
  - logical control units 476
  - logical partition 473
  - not-operational reason 386
- Channel Subsystem (CSS) 363, 472
- Channel Subsystems 475
- Channel to PCHID Assignment panel (SE) 433
  - View Cage Details 434
- Channel-to-Channel (CTC) 15, 17, 38–39
- CHPID 353, 480
- CHPID 23 and CHPID 27 136
- CHPID Details panel 150, 192, 240
- CHPID Mapping Tool (CMT) 409
- CHPID number 327, 413
- CHPID statement 50, 134, 175, 218, 413
- CHPID statements 134, 175, 218
- CHPID type
  - FC 14, 38–39
  - FCP 38–39
  - FCV 13, 42
- CHPID type FC 36
- CHPIDS 17, 38–39
- Cisco MDS 9500 Series 83
- Cisco MDS 9506 56
- Cisco MDS 9509 56
- Cisco MDS 9513 56
- Clear Subchannel (CSCH) interruption 382
- CLI command 375
- CMR Time exception 351
- CMT program 411
- CMT user's guide 411
- CNTLUNIT CUNUMBR 175, 219
- CNTLUNIT statement 134, 175, 219, 414
- CNTLUNIT statements 134, 175, 218
- Command Line Interface (CLI) 364
- command mode 10, 22, 24, 28, 34
  - CCW operation 28
  - channel program 30
- command mode open exchanges 8
- command response (CMR) 322
- command response (CMR) time 351–352, 393
- Command-line Interface 79
- Command-line interface 94
- configure a FICON point-to-point topology 120
- configuring the attached DS8000 storage subsystem 120
- CONFIGxx 350
- connecting FICON Director
  - Domain ID 54
- Context 82
- control processor (CP) 247
- Control unit 100, 377
  - logical link address 482
  - logical-path establishment 22
  - login 22
  - N\_Port address 51
  - node descriptor 35
  - node-identifier acquisition 22
  - serial number 485
  - state-change registration 22
  - type 485
- control unit (CU) 1, 9, 22, 46, 324, 350, 409, 480
- Control Unit Port (CUP) 51, 246
  - function 16
- Control unit Port (CUP)
  - optional feature) 96
- Control Unit Remote Node Identifier 483
- CPC details panel (SE) 429
  - Activation Profile 430
  - IOCDS name 430
  - product information 430
- CPU Node Descriptor 472
- CRC errors 110
- create a new switch definition 162, 204
- CSS 17, 476
- CSS image IDs 17
- CSS.CHPID
  - 1.20 150
  - 1.21 240
  - 1.22 192
- CU 4044 224–225
- CU host adapter port 0033 to FICON Director port 34 189

CU host adapter port 0133 to FICON Director #66 port 41 237  
CU host adapter port 0233 to FICON Director port 62 189  
CU host adapter port 0333 to FICON Director #66 port 11 237  
CUADD keyword 134, 136, 175, 178, 219, 226  
CUNUMBR keyword 134–136, 175–176, 178, 219–220, 226  
CUP 16, 103  
Cyclical Redundancy Checking (CRC) 53

## D

D IOS,CONFIG 141, 182, 230  
D IOS,ZHPF 147, 189, 237  
D M=CHP(20) 147  
D M=CHP(21) 237  
D M=CHP(22) 189  
D M=DEV(9000) 148, 191, 239  
D\_port 376  
Data Center Fabric Manager (DCFM) 63, 246–248, 362, 400  
DCFM server 247, 400  
    IP address 259  
    welcome window 259  
Define Access List 124, 161, 203  
Define Device to Operating System 126  
Define Device to Operating System Configuration 165, 207  
defined channel path  
    path status 486  
Delta view 369  
Dense Wavelength Division Multiplexing (DWDM) 88  
Descriptor field names 138, 180, 228  
detailed FICON channel information  
    IFCC Threshold Exceeded 437  
detailed FICON channel information (SE) 436  
    CHPID status 437  
    Definition Error 437  
    FICON CHPID and switch port link error logging chart 438  
    IFCC Threshold Exceeded 437  
    Sequence Not Permitted 438  
    Sequence Time Out 438  
Device Connectivity 372  
Device Number 353, 476, 490  
Device Parameters / Features 126, 165, 207  
Device Status (DS) 477  
device type 191, 239, 412, 476  
    cam display information 486  
DEVSERV command 476  
    other uses 486  
Devserv command 484  
Diagnostic Port Test 372, 376  
diagnostic procedure 355  
Director or Logical Switch SupportSave 365  
Director port (DP) 325  
Display Matrix command 350, 361  
Domain ID 50, 157, 199, 246, 413  
Driver level (HMC/SE) 426  
DS P 361, 484

DS8000 32, 36  
    port identifiers 417  
DS8000 storage  
    control unit 146, 188, 236  
    controller 122, 159, 201  
    system 120  
DS8000 storage controller 122  
DS8000 storage device 156  
DS8000 storage subsystem 156, 198  
DWDM 89  
Dynamic Channel Management 18  
Dynamic Channel Management (DCM) 482  
Dynamic Channel Path Management 96  
Dynamic Load Sharing (DLS) 65, 246  
Dynamic Path Selection (DPS) 65, 381  
Dynamic Pathing Switch (DPS) 378

## E

E\_Port 6, 282, 288  
E\_port 370  
E\_ports 372  
EEQINSTM 410  
Electrical Loopback Test (E-Wrap) 373  
Element management 79, 95  
Element Manager 368  
Enterprise Class (EC) 398, 473  
Enterprise Fabric Connectivity Manager (EFCM) 250  
Environmental Record, Editing, and Printing (EREP) 357  
errdump, CLI command 365, 375  
ESCON Director 51  
ESCON solution 13  
establish a logical path 377  
Establish Logical Path (ELP) 23  
Evaluation tools 111  
Exchange Based Routing 109  
Extended distance FICON 104  
Extended Link Service (ELS) 27  
Extended Link Services (ELS) 10  
Extended Status Word Timeout (ESWT) 379  
EXTENDED SUBCHANNEL LOGOUT DATA 360  
Extended Subchannel Logout Data 380  
Extended Subchannel Logout Data (ESLD) 358

## F

F\_Port 6  
F\_Port Login 22  
fabric 6  
Fabric Device Sharing 372  
Fabric login (FLOGI) 22, 36  
Fabric management 79, 95  
Fabric Manager (FM) 72  
fabric port (F\_PORT) 6, 72, 282  
Fabric security 91  
Fabric Shortest Path First 108  
FC 0409 120, 156, 199  
FC 0709 and FC 7092 122, 159, 201  
FC link 7, 24, 47  
FC Ping 372  
FC Trace Route 370

- fcping, CLI command 372, 375–376
- FC-SB-2, SB-3, SB-4 Web site 5
- FCTC control unit 15, 225, 489
- FCTC control units and device definitions 136, 177, 224
- Feature code 37, 58, 200, 403
- Fiber cabling infrastructure 120, 156, 198
- Fiber Channel Analyzer panel (HMC) 466
  - Fiber Channel Analyzer panel in service mode 468
- fiber optic cable 7, 50, 121, 198
- fiber optic cables 147, 189, 237
- Fibre Channel
  - adapter 26
  - architecture 1, 51
  - bit error 358
  - data traffic 73
  - fabric 51
  - link 26
  - logical switch 313
  - Physical 4
  - Physical and Signaling Standard 9
  - physical framing 5
  - physical interface 4
  - port 358
  - Protocol 14, 27, 55
  - Protocol traffic 38–39
  - Routing 63
  - Security Protocol 72
  - Single Byte Command Sets-3 21
  - Single Byte Command Sets-4 21
  - Single-Byte-3 9
  - Single-Byte-4 9
  - standard 4
  - Switch Fabric and Switch Control Requirements 9
- Fibre Channel (FC) 3, 21, 46, 313, 320, 357
- Fibre Channel Arbitrated Loop (FC-AL) 6
- Fibre Channel ELS Echo request 376
- FICON 1, 3, 21, 45–46, 155, 245, 319, 409, 489
- FICON advantage 13
- FICON channel 6, 22, 46, 353–354, 401, 413
  - CHPID type 413
  - different flavors 321
  - fiber links 147
  - full duplex data flow capabilities 491
  - IOS0511 message 379
  - PEND time ends 322
  - physical connectivity 49
  - physical transmission path 48
  - topologies 46
- FICON channel (FC) 327
- FICON Channel-to-Channel (FCTC) 135, 176, 220
- FICON CTC Image-ID Worksheet 405
- FICON Director 1, 14, 23, 46, 101, 156, 245–246, 321, 350, 398, 413, 481, 489
  - arbitrary number 218
  - Basic functions 50
  - control unit 47
  - cooling capability 49
  - CP cards 248
  - default buffer credit assignment 402
  - Domain ID range 50
  - fabric binding database 54
  - FC links 47
  - IBM qualification testing 56
  - in-band management 51
  - ISL port 299
  - Other settings 188, 236
  - physical port 16
  - switch module 49
  - Switch number 448
- FICON Director Configuration Worksheet 405
- FICON Directors 198
- FICON ELS (Extended Link Services) error detected error I/O Trace 455
- FICON environment 4, 45, 75, 117, 119, 197, 278, 317, 349, 360, 398, 471–472, 489
- FICON Express 42
- FICON Express features are supported on z10 servers only if carried forward on a server upgrade. 13
- FICON Express2 3, 24, 41–42
- FICON Express4 24, 36, 39, 41, 122, 158, 200
  - 10KM LX 37
  - 10KM LX feature 40
  - 4 KM LX feature 40
  - 4KM LX 37
  - CHPIDs 40
  - feature 39
  - SX 37
- FICON Express4 10KM LX 40
- FICON Express4 4KM LX 40
- FICON Express4 SX 41
- FICON Express4-2C 4 KM LX 41
- FICON Express4-2C SX 41
- FICON Express8 3, 24, 34, 36, 39, 122, 158, 200, 398
- FICON Express8 10KM LX 39
- FICON Express8 SX 39
- FICON Express8S 3, 12, 24, 30, 34, 36–37, 41, 122, 158, 200–201
- FICON Express8S 10KM LX 38
- FICON Express8S LX 120, 156, 199
- FICON Express8S SX 38
- FICON External error I/O Trace 459
- FICON FC-SB3 and SB-4 Web site 27
- FICON feature 9, 21, 200, 401
- FICON I/O request 25
- FICON icon (SE) 435
- FICON IFCC and Other errors panel (SE) 453
  - Channel Hardware Failure 461
  - FICON ELS (Extended Link Services) error detected error 455
  - FICON External error 459
  - FICON IFCC detected error 454
  - FICON timeout error 457
  - PPE (Purge Path Extended) Log 461
- FICON IFCC detected error I/O Trace 454
- FICON link incident records 357
- FICON link incident reporting 36
- FICON point-to-point configuratio 120
- Ficon problem determination 369

FICON Purge Path Extended (PPE) 358  
 FICON switched configuration 156  
 FICON switched topology 156  
 FICON timeout error  
   CCW# 458  
   CHI (Channel Image) 458  
   CMD# 458  
   CUI (Control Unit Image) 458  
   Dev# 458  
   DHF (Device Header Flag) 457  
   DIB (Device Information Block) 458  
   I/O Trace 457  
   IUI (Information Unit ID) 457  
   LA (Link Address) 458  
   Token 458  
   UA (Unit Address) 458  
 ficonshow, CLI command 376  
 FL\_Port 6  
 Forward Route 372  
 FOS 80  
 Frame pacing 103  
 Frame-based trunking 108

## G

G\_Port 6  
 Gbps 24, 49, 284, 398  
 Geographically Dispersed Parallel Sysplex™ (GDPS®)  
 84  
 graphical user interface 410  
 graphical user interface (GUI) 72, 248  
 grep, CLI operand command 376

## H

Hardware Configuration Definition (HCD) 9, 18, 50, 350,  
 409, 473  
 Hardware Configuration Manager (HCM) 9, 18, 409  
 Hardware Management  
   Console 352, 473  
 Hardware Management Console 426  
   Driver level 426  
   LIC level 426  
   log on 427  
   Password 427  
   Single Object Operations (SOO) 427  
   Userid 427  
 Hardware Management Console (HMC) 149, 192, 240  
   Tasks pad (HMC/SE) 429  
 Hardware Managemnt Console (HMC)  
   Navigation Pane (HMC/SE) 429  
   Userid 427  
 Hardware Message 354  
 Hardware System Area (HSA) 22, 123, 160, 202, 409,  
 473  
 Hwareware Management Console (HMC)  
   Password 427  
 HCD User's Guide 123, 160, 202, 410  
 HCM  
   Connect... 131, 169  
   Create --> Switch... 169, 211

CTC definition 221  
 Edit --> Controller... 131, 170, 214  
 Edit --> Processor... 129, 167, 209  
 Esoterics... 174  
 I/O Devices... 133, 174, 217  
 ISL ports 212  
   Locate --> Processor... 129, 167, 209  
 HCM User's Guide 123, 160, 202, 410  
 help, CLI command 376  
 High integrity 92  
 High Performance FICON for System z 12  
 High Performance FICON for System z (zHPF) 10, 32  
 HMC 353  
 HMC/SE FICON Link Incident log 357  
 HMC/SE problem determination 353  
 Host Bay Adapter (HBA) 354, 482  
 HOTIO 384  
 Hyper Parallel Access Volume (HyperPAV) 9  
 HyperPAV 105–106  
 HYPERPAV ALIAS 483

## I

I/O Configuration  
   Program 409, 473  
 I/O configuration 475  
 I/O Configuration Data Set (IOCDS) 352, 409  
 I/O Configuration Program (ICP IOCP) 411  
 I/O Configuration Program (IOCP) 355  
 I/O definition  
   file 179, 227, 323  
 I/O Definition (IODF) dataset 350  
 I/O definition file (IODF) 123, 160, 202, 410  
 I/O definitions 122, 159, 201  
 I/O device 16, 25, 409, 482  
   unit addresses 414  
 I/O hardware configuration 350  
 I/O information 323  
 I/O interruptions 358, 374  
 I/O Messages 352  
 I/O operation 17, 22, 46, 324, 377, 476  
   exchange pair 27  
   FICON architecture 24  
   MIH value 31  
 I/O operations 103  
 I/O path 353  
 I/O performance measurement points 321  
 I/O priority queuing (IOQ) 331  
 I/O processor (IOP) 326  
 I/O request 24, 54, 326, 383–384, 474  
   FC-FS FC-2 frame 26  
 I/O specification 123, 160, 202  
 IBM DS8000 Storage subsystem 156  
 IBM DS8000 storage subsystem 120, 198  
 IBM FICON aggregation study 111  
 IBM Health Check 393  
 IBM Health Check function for IOS 351  
 IBM Network Advisor 95, 332, 362, 375  
 IBM Network Advisor PD Tools 353, 364  
 IBM RMF Web site 321  
 IBM SAN384B-2 FICON Director 198–199

IBM SAN768B 56  
 IBM SAN768B-2 FICON Director 156–157, 198  
 IBM System Storage b-type 83  
 IBM System Storage Web site 113  
 IBM System z196 server 120  
 IBM Tivoli System Automation for z/OS 97  
 IBM zEnterprise 196 server 156  
 IBM zEnterprise z196 server 198  
 iconshow rmid, CLI command 364  
 ID statement 412  
 IEACMDxx 350  
 IECIOSxx 384, 389  
 IECIOSxx to ZHPF=YES 147, 190, 238  
 Indirect Address List (IDAL) 29  
 Information Unit 26  
 Information Unit (IU) 8, 35  
 Initialization Complete 153, 195, 243  
 in-order delivery (IOD) 65, 246  
 Input Output Configuration Program (IOCP) 9, 50  
 Input/output (I/O) 3, 21, 321  
 Input/Output architecture 16  
 Input/Output Configuration Data Set (IOCDS) 410  
 Input/Output Supervisor (IOS) 24  
 Insistent Domain (ID) 246  
 Inter-Chassis Link (ICL) 63–64  
 Interface Control Check (IFCC) 378  
 Interface Control Checks (IFCC) 354  
 InterNational Committee for Information Technology Standards (INCITS) 4  
 Interruption Response Block (IRB) 29  
 Inter-switch link 102  
 Inter-Switch Link (ISL) 47, 50, 199, 246, 278  
 Inter-switch links (ISL) 87  
 Invalid Transmission Words 110  
 IOCDS 484  
 IOCDS files 137, 179, 227  
 IOCDS List 139, 181, 229  
 IOCDS on the Support Element 137, 179, 227  
 IOCP and z/OS I/O configuration 412  
 IOCP statements and keywords used for FICON 409  
 IOCP successfully completed 140  
 IOCP User's Guide 412  
 IODEVICE Address 135, 176  
 IODEVICE statement 134, 175–176, 218, 220, 414  
 IODF data (ID) 475  
 IOS 353  
 IOS\_CMRTIME\_MONITOR 351, 393  
 IOS001E, Inoperative Path 377  
 IOS002A, no paths available 377  
 IOS050I 380  
 IOS050I, channel detected error 378  
 IOS051I 380  
 IOS051I, Interface Timeout Detected 379  
 IOS052I, Channel Event Reported 379  
 IOS054I, Errors Detected by 380  
 IOS071 380  
 IOS071I 383  
 IOS075E, Recurring MIH Condition 381  
 IOS076E 382  
 IOS077E 383

IOS100I, Device Boxed 383  
 IOS101I, Device Foced Offline or Boxed 384  
 IOS102I, Device Boxed or Forced Offline 384  
 IOS107I, Box Processing Defferred 385  
 IOS162A, Chpid Alert 385  
 IOS163A, Chpid Alert 386  
 IOS167I, Device Restored by Subchannel revocery 386  
 IOS2001I, Inoperative Paths 386  
 IOS2002I, Inoperative Paths 388  
 IOS202E, Channel Path forced Offline 389  
 IOS203I, Channel Path Successfully Recovered 389  
 IOS210I, Path Recovery Initiated for Path 389  
 IOS288A, System-Initiated Reconfiguration in progress for Chpid 391  
 IOS289I, System-initiated online/offline reconfiguration has completed 391  
 IOS580E 357, 380  
 IOS580E, link detraded reporting 392  
 IOS581E 357, 380  
 IOS581E, link degraded reporting 392  
 IOSHC112E 351, 393  
 IOSHC112E, Analysis of Command Response Time 393  
 IOSHC113I 351  
 IOSHC113I, Command Response Time Report 393  
 IP address 73, 246, 248, 400  
 IPL from an alternate subchannel set 13  
 ISL connection from FICON Director #65 port B4 to FICON Director #66 port 00 237  
 ISL connection from FICON Director #65 port C6 to FICON Director #66 port 20 237  
 ISL Trunking 72  
 ISLs 47, 102, 199, 246, 399  
 IU pacing protocol 35

## K

km transceivers 40

## L

L\_Port 6  
 LC Duplex 39  
 LC Duplex connector 39, 121, 157, 199  
 LC duplex connector 122, 159, 201  
 LIC level (HMC/SE) 426  
 Link address 6519 195  
 Link address 6534 195  
 Link address 65FE 196  
 Link address 6641 243  
 link address of 0D 153  
 Link Error Status Block (LESB) 360  
 LINK keyword 175, 219, 226  
 Link Maintenance Information (LMI) 36, 357  
 Link Maintenance Information Detail Edit (A2) Report 357  
 Link Maintenance Information Detail Edit (A2) report 380  
 Link Saturation Test 373  
 Link-Incident-Record Registration (LIRR) 22  
 LISL 102  
 LMI records 357  
 Local switching 108

- Log on (HMC/SE) 427
- logical and physical aspects 410
- logical control unit (LCU) 157, 199, 325, 401
- logical control units (9000 and 9100) 121
- logical partition
  - channel subsystem 476
  - high number 489
- logical partition (LP) 17, 67, 353, 489
- Logical Path 23, 47, 354, 443
- Logical switch
  - scalability.Partition Brocade SAN768B 65
- logical switch 65, 246, 270, 365, 401
  - management isolation 65
  - port numbering 273
- logical switch types 66
- LOGREC 357
- Long Wave (LW) 60
- Longitudinal Redundancy Checking (LRC) 53
- longwave laser 121, 198
- longwave laser (LX) 157
- LookAt, IBM online facility 377
- Loss of Sync 110
- Lossless DLS 109
- lossless DLS 278, 402
- LPAR A11 136, 177, 221
- LPAR A13 136, 177, 221
- LPAR A23 472

## M

- Managemen 94
- Master Log 364
- Measure link distance 373
- MIDAW 149, 191, 239
- MIDAWs 106
- MIF 18
- MIF ID 413
- Missing Interrupt Handler 31, 96
- Missing Interrupt Handler (MIH) 31
- MM 62.5 43
- Mode Conditioning Patch (MCP) 43
- mode conditioning patch (MCP) 40
- Modified Indirect Address List (MIDAL) 29
- Modified Indirect Data Address Word 106
- Modify Device Parameters 125, 164, 206
- Multiple allegiance 105
- Multiple Channel Subsystem 98
- Multiple Image Facility 98
- MULTIPLE IMAGE FACILITY (MIF) 472
- Multiple Image Facility (MIF) 18
- Multiple Subchannel Sets 98, 106

## N

- N\_Port 4, 6, 22
- N\_Port virtualization (NPIV) 6
- Native FICON 38–39, 327
  - average number 327
- NL\_Port 6
- node 6
- Node Descriptor 362

- node descriptor 362
- Node Descriptor (ND) 24, 261, 357, 472, 482
- node status 362
- node type 362
- Node-Element Descriptor 482
- node-identification data 362
- NOTPART keyword 134, 136, 175, 178, 219, 226
- NPIV 6

## O

- ODBC 260
- Open exchange 34
- open exchange 27
- Open Exchanges 8
- operating system 25, 51, 120, 156, 198, 479, 490
- Operating system configurations 126, 164, 206
- Operation Request Block
  - additional parameters 25
- Operation Request Block (ORB) 25
- Optical Loopback Test (O-Wrap) 373
- Optical Power Measurement panel
  - formula to convert from dBm to mW 451
  - Minimum and maximum acceptable power specifications 451
  - Optical Power Display at Brocade FICON Switch 453
- optical-to-electrical-to-optical 87
- OSCONFIG 352, 478, 484

## P

- Parallel Access Volume 105
- Parallel Access Volume (PAV) 9, 100, 330
- PARTITION keyword 226
- Password (HMC/SE) 427
- Path Available Mask (PAM) 444
- PATH keyword 134, 136, 175, 177, 219, 226, 414
- Path Not Operational Mask (PNOM) 377
- Path Operational Mask (POM) 377, 444, 482
- PATH\_SCOPE=CU 389
- PAV 100
- PCHID 398
  - Essential information 150
  - owning image 150
- PCHID 398 (CHPID 20) 120
- PCHID 398 to DS8000 host adapter port 0203 147
- PCHID 399
  - Essential information 240
  - owning image 240
- PCHID 399 (CHPID 21) 199
- PCHID 399 to FICON Director #65 port 21 237
- PCHID 39C
  - Essential information 192
  - owning image 192
- PCHID 39C CHPID 22) 156
- PCHID 39C to FICON Director port 00 189
- PCHID 3E0 (CHPID 24) 120
- PCHID 3E0 to DS8000 host adapter port 0303 147
- PCHID 3E1 (CHPID 25) 199
- PCHID 3E1 to FICON Director #65 port A3 237
- PCHID 3E4 (CHPID 26) 156

- PCHID 3E4 to FICON Director port 19 189
- PCHID detail 192, 240
- PCHID details 149
- PCHID keyword 134, 175, 219
- PDCM 91
- Performance 102
- performance monitoring 321, 332
- Permanent I/O error 384
- persistent 35
- Persistent IU pacing 36
- Physical connectivity 113
- physical control units 476
- planned and unplanned changes 351
- Platform and name registration in FICON channel 15
- Point-to-point 22, 86
- Point-to-point (direct connection) 14
- point-to-point configuration 413
- point-to-point FICON Channel-to-Channel (FCTC) 135
- point-to-point topology 119, 409
- Port 110
- port 6
- port address 9, 50, 279, 401
- port addressing 101
- Port Administration 368
- Port Based Routing 109
- Port characteristics 101
- Port Decommission 374
- Port Display 362
- Port fencing 110
- Port login (PLOGI) 22, 36
- port number 50, 199, 313
- Port number 0033 157
- port number 0233 157
- Port numbers 0203 and 0303 120
- Port Statistics 369
- port type 5, 24, 246, 282, 333
- port x00 and port x19 157
- port x34 and x62 157
- port x92 and port xA3 199
- Port-Based Routing (PBR) 246
- portcfgdport, CLI command 375–376
- Port-channel 109
- portdecom, CLI command 375–376
- portdporttest, CLI command 375–376
- porterrshow, CLI command 370, 376
- portstatsclear, CLI command 370, 376
- portstatsshow, CLI command 370, 376
- Power-on Reset (POR) 123, 160, 202
- PPE (Purge Path Extended) Log
  - Architected Error Code in FC-SB4 465
  - Channel LESB 463
  - Channels FPort LESB 463
  - CU FPort LESB 464
  - CU NPort LESB 464
  - CU RNID data 464
  - Invalid CRC Count 465
  - Invalid Transmission Word 464
  - Link failure counters 464
  - PPE Trace sample 463
  - Validity Flags 463

- PPE error code 360
- PPRC Secondary Suspend (PSEC-SUSP) 488
- pre-test cables 375
- pre-test cables and optics 375
- problem determination (PD) 147, 189, 237
- Process Login (PRLI) 10
- process login (PRLI) 32
- process logout (PRLO) 32
- Product / Host SupportSave 366
- Prohibit Dynamic Connectivity Mask 91
- Purge Path Extended (PPE) 358, 450
  - I/O Trace 462
  - PPE (Purge Path Extended) Log 461
  - Purge Path Extended Log panel 462

## Q

- Quality of Service 111
- Quality of Service (QoS) 246

## R

- RCHP - Reset Channel Path (System Reset) 384
- recovering I/O 381
- Redbooks Web site 495
  - Contact us xiv
- Registered State Change Notification (RSCN) 280
- Remote Node Identification (RNID) 441, 483
- Repository 367
- Request Node ID (RNID) 10
- Request Node Identifier (RNID) 23
- Requirements 81
- Resource Measurement Facility (RMF) 320–322
- RESOURCE statement 412
- Reverse Route 372
- RMF 96, 322
- RNID 483
- role-based access control 94
- Routing information 371

## S

- SAN Display 362
- SAN Product 367
- SAN768B-2 FICON Director 199
- SCHSET keyword 135, 176, 220
- Secure File Transfer Protocol (SFTP) 72
- serial number 362
- Server Management Console (SMC) 336
- Service Information Messages (SIMs) 96
- SETIOS ZHPF, YES 147, 237
- SETIOS ZHPF, YES. 189
- SETIOS, z/OS system command 390
- Setting the display to hex 261
- Short Wave (SW) 60
- Simple Network Management Protocol (SNMP) 49
- single mode (SM) 62, 121, 157, 198
- Single Object Operation (SOO) 149, 192, 240, 427
- SIOC (System I/O Configuration) Analyzer panel (SE) 469
  - Control Unit 470



- Link Load 470
- Node ID 470
- PCHID Control Unit 470
- PCHID Partition 470
- SLH EDIT report 380
- SLH record 360
- Small Form Factor Pluggable (SFP) 49
- Stand-Alone I/O Configuration Program (IOCP) 409
- Stand-Alone IOCP 411
- State-Change Registration (SCR) 22
- Storage Area Network (SAN) 55
- storage device 46
- Storage management initiative specification 95
- Store System Information (STSI) 472
- subchannel 475
- Subchannel Logout Handler (SLH) 358
- subchannel logout record 358
- subchannel set 17, 473
- subchannel set 2 (SS2) 13
- Subchannels 17
- SuperPing, CLI command 376
- Support Center 368
- Support Element (SE) 353, 409, 426
- support element (SE)
  - Driver level 426
  - LIC level 426
  - log on 427
  - Password 427
  - Single Object Operations (SOO) 427
  - Userid 427
- support element (SE)
  - enable service status 429
- Supportsave 376
- supportsave, CLI command 368, 376
- supportshow, CLI command 368, 376
- Suppot Element (SE)
  - enable service status (HMC/SE) 429
- switch address 16, 50, 199, 276, 413
- Switch Address (Switch @) 157
- Switch Address is the Domain ID 157
- Switch Connection Control (SCC) 247, 303
- Switch ID (Switch #) 157
- SWITCH keyword 175, 219
- Switched 14, 23, 86
- switched fabric 16
- switched FICON Channel-to-Channel (FCTC) 176
- switched togology 409
- switched topology 155
- SwitchID (Switch #) 199
- SYS#.SEEQINST 410
- System Activity Display (SAD) 319–320
- System Adapter Identifier (SAID) 417
- System Control Program (SCP) 17, 482
- System Information (SI) 472
- System Management Facility (SMF) 321
- System Storage Interoperation Center (SSIC) 80
- System z 3, 45, 398, 409, 489
  - High Performance FICON 38–39
- System z environment 9
- System z High Performance FICON (zHPF) 24

- System z management for FICON Directors 96
- System z10 9, 398, 473
- System zEnterprise 196 server 120, 156
- Systems Automation (SA) 51

## T

- Tag information 362
- TE port 72
- Technical Support Data 365
- Technical SupportSave 366
- TERMINAL 384
- terminology 4
- TEST DETECTED NO CONDITIONS WHICH WOULD RESULT IN ACTIVATE FAILURE 142, 183, 231
- TI zone 299
- Traffic isolation 109
- Traffic Isolation (TI) zone 64
- Traffic Isolation (TI) Zoning 355
- Traffic management 111
- transceiver 189, 237
- Transceiver type 85
- Transport Command Area (TCA) 12
- Transport Command Area Extension (TCAX) 13
- Transport Command Control Block (TCCB) 12, 29
- transport command word (TCW) 11
- Transport Control Word (TCW) 24, 28
- Transport Indirect Address List (TIDAL) 29
- Transport Indirect Data Address Work (TIDAW) 33
- Transport mode 24, 28, 34
- transport mode 10, 22
  - I/O operation 22
- Transport mode open exchanges 8
- Transport Status Block (TSB) 29
- Trivial file transfer protocol (TFTP) 55
- Trunk members 375
- trunk or ISL (TI) 246

## U

- U\_Port 6
- UCB Logical Path Mask (UCBLPM) 481
- unit address (UA) 26, 414
- Unit Control Block (UCB) 25, 485
- Unit Control Word (UCW) 475
- UNIT keyword 135, 176
- UNIT keyword shows 220
- UNITADD keyword 135, 176, 220
- Unrepeated distance 39–40
- Upper Level Protocol (ULP) 5
- urouteshow, CLI command 372, 376
- Userid (HMC/SE) 427

## V

- View Frame Layout panel (HMC/SE) 431
- view frame layout panel (SE)
  - I/O Cage and I/O Drawers 432
- Virtual Fabric 67, 270
- Virtual Fabrics 101
- Virtual Fabrics feature 65

Virtual ISL 65, 102  
Virtual Storage Area Network (VSAN) 72  
Virtualization and availability 97

## W

Wave Division Multiplexing (WDM) 311  
Wavelength Division Multiplexor (WDMs) 80  
Web Tools Performance Monitor 346  
World Wide Name 7  
World Wide Node  
    Name 24, 303  
World Wide Node (WWN) 269  
World Wide Node Name 15  
World Wide Node Name (WWNN) 54, 153, 195, 243  
World Wide Node\_Name (WWNN) 7  
World Wide Port  
    Name 24, 355  
World Wide Port Name 16  
World Wide Port Name (WWPN) 153, 195, 243  
World Wide Port\_Name (WWPN) 7  
Write Protect column 139, 181, 229  
WWNN 362  
WWPN 362

## X

XISL 102

## Z

z/OS discovery and auto-configuration (zDAC) 18  
z/OS Display Unit command 476  
z/OS FICON messages (IOS) 377  
z/OS Resource Measurement Facility (RMF) 112  
z/OS System Control Program (SCP) 476  
z/OS V1R13 120, 156, 198  
z10 BC 41  
z10 EC 43, 401  
    FICON connectivity 401  
z114 13, 32, 475  
z196 13, 32, 475  
z196 FICON Express8S 157  
z196 server 156, 198  
    configuration tasks 149, 192, 240  
    fiber optic cable link 153, 195, 243  
    PCHID 398 149  
    PCHID 399 240  
    PCHID 39C 192  
z196 servers 120  
zDAC 18–19, 409  
zEnterprise 114 9  
zEnterprise 114 (z114) 3, 21  
zEnterprise 196 9  
zEnterprise 196 (z196) 3, 21  
zEnterprise 196 Configuration Setup 410  
zEnterprise 196 server 156, 198  
zEnterprise z196 server 120  
ZHPF 149, 191, 239  
zHPF 10–11, 28, 30, 107, 112  
zHPF bi-directional (BiDi) capability 13

zHPF DB2 list PREFETCH 12  
zHPF feature 121, 157, 199, 402  
zHPF now for QSAM, BPAM, and BSAM access methods 12  
zHPF parameter in the SYS1.PARMLIB 147, 190, 238  
zHPF protocol 156, 198  
zHPF Temporary Logout 32  
Zone Config 294  
Zoning 93  
zoning 355



Redbooks

# FIGON Planning and Implementation Guide

(0.5" spine)  
0.475" x 0.873"  
250 x 459 pages







# FICON Planning and Implementation Guide



**Learn about topologies, concepts, and terminology**

**Receive guidance for planning, implementation, and migration**

**Get realistic examples and scenarios**

This IBM Redbooks publication covers the planning, implementation, and management of IBM System z FICON environments. It discusses the FICON and Fibre Channel architectures, terminology, and supported topologies.

The book focuses on the hardware installation and the software definitions that are needed to provide connectivity to FICON environments. You will find configuration examples required to support FICON control units, FICON Channel-to-Channel (FCTC), and FICON Directors. It also discusses utilities and commands that are useful for monitoring and managing the FICON environment.

The target audience for this document includes IT Architects, data center planners, SAN administrators, and system programmers who plan for and configure FICON environments. You are expected to have a basic understanding of IBM zEnterprise and System z10 hardware, HCD, HCM or IOCP, and a broad understanding of the Fibre Channel and FICON architectures.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

## **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)