




System z10 Business Class System Overview

SA22-1085-04

Level 04a, October 2010





System z10 Business Class System Overview

SA22-1085-04

Level 04a, October 2010

Note

Before using this information and the product it supports, read the information in “Safety” on page xiii, Appendix D, “Notices,” on page 143, and *IBM Systems Environmental Notices and User Guide*, Z125-5823.

- | This edition, SA22-1085-04, applies to the IBM System z10 Business Class (z10 BC) server. This edition replaces
- | SA22-1085-03. Technical changes to the text are indicated by a vertical bar (|) to the left of the change.

There might be a newer version of this document in a PDF file available on **Resource Link**. Go to <http://www.ibm.com/servers/resourcelink> and click **Library** on the navigation bar. A newer version is indicated by a lowercase, alphabetic letter following the form number suffix (for example: 00a, 00b, 01a, 01b).

© Copyright IBM Corporation 2008, 2010.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	ix
Tables	xi
Safety	xiii
Safety notices	xiii
World trade safety information	xiii
Laser safety information	xiii
Laser compliance	xiii
About this publication	xv
What is included in this publication.	xv
Revisions	xv
Prerequisite publications	xv
Related publications	xv
Parallel sysplex publications	xvi
OSA publications	xvi
Cryptographic publications	xvii
Miscellaneous publications	xvii
Related websites	xvii
Additional online information	xviii
Engineering change (EC) level considerations	xviii
How to send your comments	xviii
Summary of changes	xix
Chapter 1. Introduction	1
z10 BC highlights	2
z10 BC model.	6
Performance	7
Resource Link	7
Fiber optic cabling	7
z/Architecture	8
Upgrade progression	9
Unsupported features/functions	9
Chapter 2. Hardware characteristics	11
System frame configuration	11
Central Processor Complex (CPC) drawer	12
I/O drawer	18
Support Element	20
System power supply	21
Internal Battery Feature (IBF)	21
Internet Protocol Version 6	22
Multiple Subchannel Sets (MSS)	22
LPAR mode	22
Processor units.	23
Storage	23
Channels	23
LPAR time offset support	24
Server Time Protocol (STP)	25
Hardware Management Console (HMC).	25
Bolt-down kit.	25

Power sequence control	26
Additional features/functions supported	26
Monitoring and estimating CPC power consumption and temperature	26
Preplanning and setting up the Storage Area Network (SAN) environment	26
Chapter 3. Software support	27
z/OS.	28
z/VM	29
z/VSE	29
Linux on System z	30
TPF	30
Chapter 4. Channel subsystem structure	31
IOCP channel, link, and adapter definitions	32
Coupling link peer channels	33
Subchannel connectivity	34
Guidelines for maximum availability	34
InfiniBand and ICB-4 configurations	36
Planning for channel subsystem	37
PCHID assignments	38
AID assignments	39
PCHID report	39
CHPID Mapping Tool.	40
Multiple Image Facility (MIF)	40
Spanned channels	41
Internal coupling and HiperSockets channels	41
IOCP considerations	42
LPAR definition	42
Channel path definition	42
I/O device definition	43
Hardware Configuration Definition (HCD) considerations.	43
Chapter 5. I/O connectivity	45
FICON and FCP channels.	45
FICON Express8 features	45
FICON Express4 features	46
FICON Express2 features	47
FICON Express features	47
Name server registration	48
High Performance FICON for System z (zHPF)	48
Improved performance at extended distance	48
The MIDAW facility	48
Multipath Initial Program Load (IPL)	48
Purge path extended.	48
Fibre channel analysis	49
Fibre Channel Protocol (FCP) for SCSI devices	49
ESCON channels	54
ESCON converter operation	54
I/O operations control	56
I/O interface protocol.	59
ESCON channel performance	61
OSA channels	64
Supported CHPID types	64
OSA/SF	65
OSA-Express3 features.	65
OSA-Express2 features.	67

OSA-Express3 and OSA-Express2 supported functions	68
HiperSockets	70
Broadcast support.	71
IPv6 support.	71
VLAN support	71
HiperSockets Network Concentrator	71
HiperSockets Network Traffic Analyzer	71
Layer 2 (Link Layer) support	72
Multiple Write facility	73
Chapter 6. Sysplex functions	75
Parallel Sysplex	75
Parallel Sysplex coupling link connectivity	76
Peer mode	77
ISC-3 links	77
ICB-4 links	78
InfiniBand (IFB) links	78
IC links.	80
Timing-only links	80
Coupling facility.	80
CFCC considerations	81
Coupling connection considerations	83
I/O configuration considerations.	83
Server Time Protocol (STP)	83
System-managed CF structure duplexing	86
Geographically Dispersed Parallel Sysplex (GDPS)	87
GDPS/PPRC	87
GDPS/XRC	88
GDPS/Global Mirror	88
Intelligent Resource Director (IRD).	88
LPAR CPU management (clustering)	89
I/O priority queuing (IOPQ)	89
Dynamic channel path management (DCM)	90
Workload manager (WLM).	90
EAL5 certification	91
Chapter 7. Cryptography.	93
CP Assist for Cryptographic Function (CPACF)	93
Protected key CPACF	94
Enablement and disablement of DEA key and AES key functions	94
Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P	94
Crypto Express2 and Crypto Express2-1P	95
Crypto Express3 and Crypto Express3-1P	95
User-defined extensions	96
Trusted Key Entry (TKE)	96
Trusted Key Entry (TKE) with Smart Card Readers	97
Wizard for migrating cryptographic configuration data.	97
RMF monitoring	98
FIPS certification	98
Remote loading of ATM and POS keys	98
Chapter 8. Cabling	99
Fiber Quick Connect (FQC) for ESCON and FICON LX cabling	100
Cabling responsibilities	100
Cable ordering	101

Cabling report	102
Chapter 9. Hardware Management Console and Support Element	105
Hardware Management Console Application (HWMCA)	106
Hardware Management Console and Support Element enhancements for z10	
BC	107
New and revised tasks	108
Hardware Management Console and Support Element network connection	108
Hardware Management Console (HMC) features and functions.	108
Customization of the HMC or SE.	108
Status reporting	108
Service Required state	109
Degrade indicator	109
Hardware messages	109
Operating system messages	110
Problem analysis and reporting	110
Virtual RETAIN	111
Licensed Internal Code (LIC)	111
Remote I/O configuration and IOCDS management	111
Scheduled operations	111
Remote Support Facility (RSF)	112
Automation and API support	112
CPC activation	113
NTP client/server support on the HMC	113
z/VM integrated systems management.	113
Installation support for z/VM using the HMC.	114
User authentication	114
Network protocols	114
Customizable console date and time	114
System I/O configuration analyzer (SIOA).	114
Network analysis tool for Support Element communications	115
Instant messaging facility.	115
Screen capture function	115
Call-home servers selection.	115
User interface	115
User authority	116
Available media	116
Security considerations	116
Change management considerations	117
Remote operations and remote access	117
Remote manual operations	118
Remote automated operations	119
Chapter 10. Reliability, Availability, and Serviceability (RAS)	121
Reliability	121
Availability	121
Redundant I/O interconnect.	121
Plan ahead features	121
Enhanced driver maintenance	122
Program directed re-IPL	122
Concurrent MBA fanout card and HCA fanout card hot-plug	122
Processor unit (PU) sparing.	122
Processor design	123
Support Element (SE)	123
Hardware Management Console	123
Attaching to IBM service through the Internet	123

Hardware Management Console monitor system events	123
SAPs	123
Application preservation	124
Dynamic coupling facility dispatching	124
Error Correction Code (ECC)	124
Dynamic memory sparing	124
Memory scrubbing	124
Fixed HSA	124
Dynamic additions to a channel subsystem and LPARs	124
LPAR dynamic storage reconfiguration	124
CICS subsystem storage protect	125
Partial memory restart	125
Dynamic I/O configuration	125
FICON cascaded directors	125
FCP full-fabric connectivity	125
Concurrent channel upgrade	125
Dual power feeds	126
Redundant power and thermal subsystems	126
Oscillator (OSC) and External Time Reference (ETR) cards	126
Preferred Time Server and Backup Time Server	126
Concurrent hardware maintenance	127
Concurrent Licensed Internal Code (LIC) patch	127
Electronic Service Agent (Service Director)	127
Internal Battery Feature (IBF)	128
Redundant coupling links	128
Large page support	128
Customer Initiated Upgrade (CIU)	128
Capacity Upgrade on Demand (CUoD)	128
On/Off Capacity on Demand (On/Off CoD)	128
Capacity Backup (CBU)	128
Capacity for Planned Events (CPE)	129
Capacity provisioning	129
System-managed CF structure duplexing (CF duplexing)	129
Geographically Dispersed Parallel Sysplex (GDPS)	129
Concurrent undo CBU	129
Fiber optic cabling	130
CHPID Mapping Tool	130
Multipath initial program load	130
Point-to-point SMP network	130
System-initiated CHPID reconfiguration	130
Link aggregation support	130
Network Traffic Analyzer Trace facility	131
QDIO diagnostic synchronization	131
FICON purge path extended	131
FICON Express8 and FICON Express4 pluggable optics for individual servicing	131
CICS subspace group facility	132
Dynamic channel path management	132
Serviceability	132
Appendix A. System z10 BC Version 2.10.2 purpose and description	133
Preventative Service Planning (PSP) bucket considerations	133
Software corequisites	133
Engineering change (EC) considerations	133
CPC EC N24409 + MCLs	133
HMC EC N24415 + MCLs	133

Miscellaneous lower level ECs included in Version 2.10.2.	134
Appendix B. Resource Link	135
Resource Link Functions	135
Appendix C. Capacity upgrades	137
Permanent upgrades	137
Temporary upgrades	137
On/Off Capacity on Demand (On/Off CoD)	138
Capacity Backup (CBU)	139
Capacity for Planned Events (CPE)	140
Concurrent PU conversions	140
Reserved CP support in LPAR mode	140
Nondisruptive upgrades	140
Processor capacity downgrades	141
Appendix D. Notices	143
Trademarks.	144
Electronic emission notices	145
Glossary	149
Index	163

Figures

1. System z10 BC	1
2. z10 BC frame configuration	11
3. z10 BC CPC drawer	12
4. I/O drawer layout for the z10 BC	35
5. InfiniBand, ICB-4, and ISC-3 connections	37
6. Control Unit (CU) priority on ESCON channels attached to a 9034 ES connection converter	60
7. Coupling link connectivity	77
8. Cabling section of the PCHID report sample	102
9. Hardware Management Console configuration	106
10. Remote operation example configuration.	118

Tables

1. Summary of changes	xix
2. System z10 BC model structure.	7
3. PUs per z10 BC model	13
4. Channels, links, ports, and adapters summary per system	19
5. Supported operating systems for z10 BC	27
6. Channel, port, adapter maximums	31
7. Channels, links, and adapters with CHPID type	32
8. PCHID assignments for MBA fanout cards	38
9. PCHIDs assignments for I/O drawer.	38
10. AID assignments for HCA fanout cards.	39
11. Coupling link options	76
12. IOPQ in a single-system environment	90
13. Channel card feature codes and associated connector types and cable types	101
14. Software corequisites.	133
15. ECs included in Version 2.10.2	134

Safety

Safety notices

Safety notices may be printed throughout this guide. **DANGER** notices warn you of conditions or procedures that can result in death or severe personal injury.

CAUTION notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous. **Attention** notices warn you of conditions or procedures that can cause damage to machines, equipment, or programs.

There are no **DANGER** notices in this guide.

World trade safety information

Several countries require the safety information contained in product publications to be presented in their translation. If this requirement applies to your country, a safety information booklet is included in the publications package shipped with the product. The booklet contains the translated safety information with references to the US English source. Before using a US English publication to install, operate, or service this IBM® product, you must first become familiar with the related safety information in the *Systems Safety Notices*, G229-9054. You should also refer to the booklet any time you do not clearly understand any safety information in the US English publications.

Laser safety information

All System z® models can use I/O cards such as PCI adapters, ESCON®, FICON®, Open Systems Adapter (OSA), InterSystem Coupling-3 (ISC-3), or other I/O features which are fiber optic based and utilize lasers or LEDs.

Laser compliance

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION:

Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)

CAUTION:

This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

About this publication

This publication describes the design, components, functions, features, and capabilities of the IBM System z10[®] Business Class (z10[™] BC) model. It is intended for executives, data processing managers, data processing technical staff, consultants, and vendors who wish to exploit z10 BC advantages.

You should be familiar with the various publications listed in “Prerequisite publications” and “Related publications.” A glossary and an index are provided at the back of this publication.

What is included in this publication

This publication contains the following chapters and appendices:

- Chapter 1, “Introduction,” on page 1
- Chapter 2, “Hardware characteristics,” on page 11
- Chapter 3, “Software support,” on page 27
- Chapter 4, “Channel subsystem structure,” on page 31
- Chapter 5, “I/O connectivity,” on page 45
- Chapter 6, “Sysplex functions,” on page 75
- Chapter 7, “Cryptography,” on page 93
- Chapter 8, “Cabling,” on page 99
- Chapter 9, “Hardware Management Console and Support Element,” on page 105
- Chapter 10, “Reliability, Availability, and Serviceability (RAS),” on page 121
- Appendix A, “System z10 BC Version 2.10.2 purpose and description,” on page 133
- Appendix B, “Resource Link,” on page 135
- Appendix C, “Capacity upgrades,” on page 137
- Appendix D, “Notices,” on page 143.

Revisions

A technical change to the text is indicated by a vertical line to the left of the change.

Prerequisite publications

Before reading this publication you should be familiar with IBM z/Architecture[®], IBM S/390[®], and IBM Enterprise Systems Architecture/390 (ESA/390) as described in the following publications:

- *z/Architecture Principles of Operation*, SA22-7832
- *Enterprise System Architecture/390 Principles of Operation*, SA22-7201.

Related publications

Important

Please ensure that you are using the most recent version of all related documentation.

Other IBM publications that you will find helpful and that you should use along with this publication are in the following list. You can access these books from *Resource Link*[®] under the **Library** section.

- *System z Application Programming Interfaces*, SB10-7030

- *System z Common Information Model (CIM) Management Interface*, SB10-7154
- *System z10 Capacity on Demand User's Guide*, SC28-6871
- *System z CHPID Mapping Tool User's Guide*, GC28-6825
- *System z ESCON and FICON Channel-to-Channel Reference*, SB10-7034
- *System z Hardware Management Console Operations Guide*, SC28-6881
- *System z10 and z9 Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152
- *System z Input/Output Configuration Program User's Guide for ICP IOCP*, SB10-7037
- *System z10 Business Class Installation Manual*, GC28-6874
- *System z10 Business Class Installation Manual for Physical Planning*, GC28-6875
- *System z10 Processor Resource/Systems Manager Planning Guide*, SB10-7153
- *System z10 Support Element Operations Guide*, SC28-6882
- *System z Small Computer Systems (SCSI) IPL - Machine Loader Messages*, SC28-6839
- *Enterprise Systems Architecture/390 Common I/O-Device Commands and Self Description*, SA22-7204
- *Enterprise Systems Architecture/390 Data Compression*, SA22-7208
- *Enterprise Systems Architecture/390 ESCON Channel-to-Channel Adapter*, SA22-7203
- *Enterprise Systems Architecture/390 ESCON I/O Interface*, SA22-7202
- *Enterprise System Architecture/390 System 360 and System 370 I/O Interface Channel to Control Unit OEMI*, GA22-6974
- *Cabling System Optical Fiber Planning and Installation Guide*, GA27-3943
- *Introducing Enterprise Systems Connection*, GA23-0383
- *System z Planning for Fiber Optic Links*, GA23-0367
- *Storage Subsystem Library Introduction to Nonsynchronous Direct Access Storage Subsystems*, GC26-4519.

Parallel sysplex publications

A **Parallel Sysplex**® system consists of two or more z/OS® images coupled by coupling links to a common Coupling Facility and synchronized by a common time source, such as Server Time Protocol (STP) or a Sysplex Timer®. A Parallel Sysplex can be used to present a single image to the user. A Parallel Sysplex can use the coupling facility to provide data sharing among the systems participating in the Parallel Sysplex.

The following publications provide additional information to help you understand and prepare for a Parallel Sysplex that uses coupling facility for data sharing purposes.

- *z/OS Parallel Sysplex Application Migration*, SA22-7662
- *z/OS Parallel Sysplex Overview: Introducing Data Sharing and Parallelism in a Sysplex*, SA22-7661
- *z/OS MVS Setting Up a Sysplex*, SA22-7625

OSA publications

The following publication provides additional information for planning and using the OSA-Express features:

- *System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*, SA22-7935
- *System z9 and eServer zSeries z890 and z990 Open Systems Adapter-Express Integrated Console Controller User's Guide*, SA22-7990.

Cryptographic publications

The following publications provide additional information about the Cryptographic function:

- *z/OS ICSF Trusted Key Entry PCIX Workstation User's Guide*, SA23-2211
- *z/OS Integrated Cryptographic Service Facility Administrator's Guide*, SA22-7521
- *z/OS Integrated Cryptographic Service Facility Application Programmer's Guide*, SA22-7522
- *z/OS Integrated Cryptographic Service Facility Messages*, SA22-7523
- *z/OS Integrated Cryptographic Service Facility Overview*, SA22-7519
- *z/OS Integrated Cryptographic Service Facility System Programmer's Guide*, SA22-7520
- *OS/390® ICSF Trusted Key Entry Workstation User's Guide*, GA22-7430
- *OS/390 Integrated Cryptographic Service Facility Administrator's Guide*, SC23-3975
- *OS/390 Integrated Cryptographic Service Facility Application Programmer's Guide*, SC23-3976
- *OS/390 Integrated Cryptographic Service Facility Messages*, SC23-3977
- *OS/390 Integrated Cryptographic Service Facility Overview*, GC23-3972
- *OS/390 Integrated Cryptographic Service Facility System Programmer's Guide*, SC23-3974
- *Cryptographic Subsystem Concepts and Facilities*, GC22-9063
- *Data Security Through Cryptography*, GC22-9062.

Miscellaneous publications

The following publications provide additional miscellaneous information:

- *IBM 3390 Direct Access Storage Introduction*, GC26-4573
- *IBM Enterprise Storage Server Host Systems Attachment Guide*, SC26-7296
- *IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10 and E20,,* GC26-7294
- *z/OS Communications Server New Function Summary*, GC31-8771
- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Implementation Guide*, SG24-7281
- *Getting Started with InfiniBand on System z10 and System z9*, SG24-7539

Related websites

The following websites provide additional System z10 information:

Resource Link

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

Resource Link is a key element in supporting the System z10 product life cycle. Some of the main areas include:

- *Planning*
- *Education*

- *Library*
- *Services*
- *CHPID Mapping Tool*
- *Customer Initiated Upgrade (CIU)*

Supported operating systems information:

<http://www.ibm.com/systems/z/os/>

Parallel Sysplex and coupling facility information

<http://www.ibm.com/systems/z/psa/>

FICON information:

<http://www.ibm.com/systems/z/hardware/connectivity>

Linux on System z information:

- <http://www.ibm.com/systems/z/os/linux/>
- <http://www.ibm.com/developerworks/linux/linux390/>

Note: When searching, specify “Linux” instead of “All of dW.”

Additional online information

Online information is available about the z10 BC system to define tasks and to aid in completing tasks. The following information is available under the Library category on the Hardware Management Console Welcome screen:

- Application Programming Interfaces (API)
- Programming Interfaces for Java
- Common Information Model (CIM) Management Interface
- Coupling Facility Control Code (CFCC) commands
- Coupling Facility Control Code (CFCC) messages
- Hardware Management Console Operations Guide
- Support Element Operations Guide.

Help is available for panels, panel options, and fields on panels.

Engineering change (EC) level considerations

Future enhancements available for z10 BC models may be dependent on the EC level of the Central Processor Complex (CPC) and/or Hardware Management Console. Additionally, some enhancements may further be dependent on the Microcode Load (MCL) level of the EC on the CPC and/or Hardware Management Console. The required MCL level will be available to the IBM field representative.

EC levels can be tracked by accessing Resource Link, <http://www.ibm.com/servers/resourcelink>, go to **Tools -> Machine Information**.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. Send your comments by using Resource Link at <http://www.ibm.com/servers/resourcelink>. Select **Feedback** on the navigation bar on the left. You can also send an email to reslink@us.ibm.com. Be sure to include the name of the book, the form number of the book, the version of the book, if applicable, and the specific location of the text you are commenting on (for example, a page number, table number, or a heading).

Summary of changes

Summary of changes for the *System z10 BC System Overview*, SA22-1085.

Table 1. Summary of changes

Release Level	Date	Changes in Level
04a	10/2010	This revision contains editorial changes.
04	07/2010	This revision contains editorial changes.
03a	01/2010	This revision contains editorial changes and the following technical changes: <ul style="list-style-type: none"> • Network traffic analyzer support for HiperSockets™
03	10/2009	This revision contains editorial changes and the following technical changes: <ul style="list-style-type: none"> • Crypto Express3 and Crypto Express3-1P features • Cryptographic migration wizard on TKE for Crypto Express2 and Crypto Express3 • TKE 6.0 feature • TKE workstation (FC 0840) will be available 1Q2010 • For OSA-Express3 1000BASE-T Ethernet feature (FC 3367) and OSA-Express3-2P 1000BASE-T Ethernet feature (FC 3369), the ability to use both ports on a single CHPID when configured as CHPID type OSC – exploitation of four ports for FC 3367 and two ports for FC 3369 (available 1Q2010) • Optimized latency mode of OSA-Express3 operating in QDIO mode (CHPID type OSD). • Capacity for Planned Event (CPE) allows you to select the capacity to meet your business needs rather than providing temporary access to all dormant capacity (available 4Q2009). • Ability to perform problem analysis for FICON channel link errors of attached Support Elements using the Fibre channel analysis task on the HMC.
02	07/2009	This revision contains editorial changes and the following technical changes: <ul style="list-style-type: none"> • Addition of FICON Express8 feature • EAL5 certification • Linux on System z: Novell SUSE SLES 11 support.
01	04/2009	This revision contains editorial changes and the following technical changes: <ul style="list-style-type: none"> • STP enhancements • User ID and password required to download internal code from IBM Support System to removable media • CBU expiration date and test activation changes • Worldwide Port Name (WWPN) prediction tool • Increased memory size to a maximum of 248 GB.

Chapter 1. Introduction

The IBM System z10 Business Class (z10 BC) represents the continuation of the IBM System z10 platform introduced with the IBM System z10 Enterprise Class (z10 EC). The z10 BC is designed and optimized for On Demand Business. This new generation of IBM servers provides an advanced combination of reliability, availability, security, scalability, and virtualization features, together with the ability to reallocate processing power. The z10 BC is designed to match changing business priorities on demand.

The newest member of the IBM System z family delivers improved granularity and enriched functions over the System z9® BC model providing a resilient infrastructure designed to satisfy the requirements of On Demand Business.



Figure 1. System z10 BC

The z10 BC and its wide range of capacity settings is designed to provide:

- Increased processing power
- Increased memory
- Improved availability
- Increased I/O capacity
- Improved security
- Enhanced network and On Demand offerings
- Enhanced system management
- Enhanced virtualization.

The z10 BC allows virtualization of resources such as:

- Sharing without boundaries.
- Empowerment of business by applying intelligence to adapt and optimize to changing requirements.

- Smart and secure management of global transactions.
- Positioning the mainframe at the center of a heterogeneous on-demand infrastructure.

To address the growing complexity of fiber optic connectivity in the Information Technology (IT) infrastructure, IBM Site and Facilities Services offers scalable fiber optic cabling services to satisfy e-business infrastructure requirements at both the product-level and the enterprise-level. Refer to Chapter 8, “Cabling,” on page 99 for more information. You can also access Resource Link at <http://www.ibm.com/servers/resourcelink> and click **Services** on the navigation bar for the network cabling services.

z10 BC highlights

The z10 BC is designed and optimized for On Demand business. The z10 BC provides:

- **One hardware model (E10)**
- **Up to 10 customer processor units (PUs)**
- **Up to five processor units (PUs) available for subcapacity use and 130 available subcapacity settings**
- **Up to 248 Gigabytes available real memory per server.**

- **Coupling using InfiniBand®**

A System z10 to System z10 connection is provided by an InfiniBand Double Data Rate (12x IB-DDR) fiber optical link with a link data rate of 6 Gigabytes per second (GBps) and maximum link distance of 150 meters (492 feet), or by an 1x IB-SDR fiber optical link or 1x IB-DDR fiber optical link with a maximum unrepeated distance of 10 kilometers (6.2 miles) and a maximum repeated distance of 100 kilometers (62 miles). 1x IB-SDR has a link data rate of 2.5 Gbps (Gigabits). 1x IB-DDR has a link data rate of 5 Gbps (Gigabits).

A System z10 to System z9 connection is provided by an InfiniBand Single Data Rate (12x IB-SDR) fiber optical link with a link data rate of 3 Gigabytes per second (GBps) and maximum link distance of 150 meters (492 feet).

- **Integrated Cluster Bus-4 (ICB-4)**

The ICB-4 link is a member of the family of Coupling Link options. ICB-4 operates at STI speeds, and it is used by coupled servers to pass information back and forth over high speed links in a Parallel Sysplex environment when the distance between servers is no greater than 7 meters or 23 feet.

- **InterSystem Coupling-3 (ISC-3)**

The ISC-3 link is a member of the family of coupling link options. ISC-3 has a link data rate of 2 Gbps. It is used by coupled servers to pass information back and forth over high speed links in a Parallel Sysplex environment for unrepeated distances up to 10 km (6 miles) and repeated distances up to 100 km (62 miles).

- **IBM System z10 Integrated Information Processor (zIIP)**

zIIP is a specialty engine designed to help improve resource optimization and lower the cost of eligible workloads enhancing the role of the mainframe as the data hub of the enterprise.

- **IBM System z10 Application Assist Processor (zAAP)**

zAAP is a specialized processor unit that provides an attractively priced Java execution environment if you desire the traditional qualities of service and the integration advantages of the System z10 BC platform.

- **Integrated Facility for Linux (IFL)**

An IFL is a specialty engine that provides additional processing capacity exclusively for Linux on System z workloads.

- **Internal Coupling Facility (ICF)**

An ICF is a specialty engine that provides additional processing capability exclusively for the execution of the Coupling Facility Control Code (CFCC) in a coupling facility partition.

- **MBA fanout card**

Memory Bus Adapter (MBA) fanout cards contain the STI connectors that allow you to connect the ICB-4 cable from one machine to another. This is typically used when connecting a z10 BC server to a z9® EC or z9 BC server.

- **HCA (Host Channel Adapter) fanout cards (HCA2-C (copper), HCA2-O (optical), and HCA2-O LR (optical))**

HCA2-C (copper) cards have InfiniBand connections used for internal I/O on a z10 BC. HCA2-O and HCA2-O LR (optical) cards have InfiniBand connections used for coupling on a z10 BC.

- **8 GB (Gigabytes) fixed size Hardware System Area (HSA)**

- **Server Time Protocol feature provides:**

- Improved time synchronization for z10 BC, z10 EC, z9 BC, z9 EC, z890, and z990 servers
- Increased multisite sysplex distance to 100 km
- Coexistence of servers and coupling facilities (CFs) synchronized in an ETR network with servers and CFs that are synchronized with Coordinated Server Time (CST)
- Concurrent migration from an ETR network
- Messaging over ICB-4 links, ISC-3 links, and InfiniBand (IFB) links
- NTP client support. The NTP client attaches to an NTP server that will provide time accuracy across heterogeneous platforms in an enterprise.
- Enhanced accuracy to an external time source utilizing pulse per second (PPS) output from an NTP server
- Use of the HMC as an NTP server configured for use as the external time source
- Continuous availability of NTP servers used as an external time source
- Enhanced STP recover when the Internal Battery Feature is in use
- Ability to save the STP configuration and time information across Power on Resets (POR) or power outages for a single or dual server STP-only CTN
- Automation of STP CTN reconfiguration using the System z Application Programming Interface (API)
- Ability to notify z/OS when events related to accessing an external time source occur.

- **Internet Protocol Version 6 (IPv6) support**

IPv6 is available for the HMC and SE customer network, the TKE network connection to operating system images, and HiperSockets. IPv6 is the protocol designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4. IPv6 expands the IP address space from 32 bits to 128 bits enabling a far greater number of unique IP addresses.

- **Power estimating and monitoring functions:**

- Power estimator tool on Resource Link
- Monitoring of power consumption and thermal loading on the Activity task on the HMC
- Support for IBM Systems Director Active Energy Manager™ for x86, IBM Systems Director Active Energy Manager for POWER®, and IBM Systems

Director Active Energy Manager for Linux on System z, which can monitor power and thermal data for z10 BC, as well as other systems.

- **Capacity on Demand functions**
 - Ability to perform a permanent LICCC upgrade while temporary resource are active
 - Ability to install and activate multiple temporary records at any given time
 - Ability to activate partial resources on a single temporary record
 - Disaster recovery solutions:
 - Capacity for Planned Events (CPE) - Short range - 3 days
 - Capacity Backup (CBU) - Long range - 90 days
 - Capacity provisioning, which provides a means of managing your processing capacity based on business needs
 - Ability to prepay for On/Off CoD upgrades
 - Ability to set spending limits when ordering an On/Off CoD record
 - 130 available subcapacity settings.
- **HiperSockets** performance improvements with the Multiple Write facility support and ability to host both non-IP workloads (with Layer 2 support) and IP workloads (with Layer 3 support). Ability to diagnose network problems using a network traffic analysis tool.
- **Large page support (1 megabyte pages)** provides performance improvement for a select set of applications, primarily long running memory access intensive applications.
- **Reduced impact of planned and unplanned server outages** through:
 - Redundant I/O interconnect
 - Enhanced driver maintenance
 - Dynamic oscillator switchover
 - Program directed re-IPL
 - Concurrent MBA fanout card hot-plug
 - System-initiated CHPID reconfiguration
 - Concurrent HCA fanout card hot-plug and rebalance.
- **Enhanced driver maintenance** allows Licensed Internal Code (LIC) updates to be performed in support of new features and functions. When properly configured, the z10 BC is designed to support activating a selected new LIC level concurrently. Certain LIC updates will not be supported by this function.
- **Redundant I/O interconnect** helps maintain critical connections to devices. The z10 BC is designed so that access to I/O is not lost in the event of a failure in an HCA fanout card, IFB cable or the subsequent repair.
- **Up to 30 logical partitions (LPARs)**
- **Server consolidation**

The expanded capacity and enhancements to the I/O infrastructure facilitates the consolidation of multiple servers into one z10 BC with increased memory, and additional I/O, which may allow you to reduce the number of servers while hosting additional applications.

The z10 BC provides the ability to define up to two logical channel subsystems (LCSS). Each LCSS is capable of supporting up to 256 CHPID definitions and 30 LPARs.
- **Non-raised floor support**

- **Frame bolt-down kit**

A bolt-down kit is available for a low raised floor installation (9 to 13 inches) and a high raised floor installation (12 to 22 inches).

- **ESCON (16 ports)**

- **FICON Express8, FICON Express4, FICON Express2, and FICON Express**

FICON Express8 features:

- FICON Express8 10KM LX (4 channels per feature)
- FICON Express8 SX (4 channels per feature)

FICON Express4 features:

- FICON Express4 10KM LX (4 channels per feature)
- FICON Express4 4KM LX (4 channels per feature)
- FICON Express4 SX (4 channels per feature)
- FICON Express4-2C 4KM LX (2 channels per feature)
- FICON Express4-2C 4KM SX (2 channels per feature)

FICON Express2 features:

- FICON Express2 LX (4 channels per feature)
- FICON Express2 SX (4 channels per feature)

FICON Express features:

- FICON Express LX (2 channels per feature)
- FICON Express SX (2 channels per feature)

Enhancements:

- zHPF (FICON Express8, FICON Express4 and FICON Express2 features (CHPID type FC))
- Extended distance (FICON Express8, FICON Express4 and FICON Express2 features (CHPID type FC))
- SCSI IPL function is part of the base system.

- **OSA-Express3 and OSA-Express2**

NOTE: OSA-Express features are no longer supported.

OSA-Express3 features:

- OSA-Express3 GbE LX (4 ports per feature)
- OSA-Express3 GbE SX (4 ports per feature)
- OSA-Express3 1000BASE-T Ethernet (4 ports per feature)
- OSA-Express3-2P 1000BASE-T Ethernet (2 ports per feature)
- OSA-Express3 10 GbE LR (2 ports per feature)
- OSA-Express3 10 GbE SR (2 ports per feature)
- OSA-Express3-2P GbE SX (2 ports per feature)

OSA-Express2 features:

- OSA-Express2 GbE LX (2 ports per feature)
- OSA-Express2 GbE SX (2 ports per feature)
- OSA-Express2 1000BASE-T Ethernet (2 ports per feature)
- OSA-Express2 10 GbE LR (1 port per feature)

- **Cryptographic options:**

- Configurable Crypto Express3 and Crypto Express3-1P features.
- Configurable Crypto Express2 and Crypto Express2-1P features. (Crypto Express2 replaces PCICA and PCIXCC.)
- CP Assist for Cryptographic Function (CPACF), which delivers cryptographic support on every PU with data encryption/decryption. CPACF also provides a high performance secure key function that ensures the privacy of key material used for encryption operations.

CPACF support includes AES for 128-, 192- and 256-bit keys; SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 for message digest; PRNG, DES, and TDES

- Using the Support Element, you can enable or disable the encrypt DEA key and encrypt AES key functions of the CPACF.
- User Defined Extension (UDX), which is supported by Crypto Express3, Crypto Express3-1P, Crypto Express2, and Crypto Express2-1P.
 - Remote loading of ATMs and POS keys.
 - Dynamically add, move, or delete a Crypto Express3, Crypto Express3-1P, Crypto Express2, or Crypto Express2-1P feature to or from an LPAR.
 - Cryptographic migration wizard on TKE for migrating configuration data from one Cryptographic coprocessor to another Cryptographic coprocessor.
 - The tamper-resistant hardware security module, which is contained within the Crypto Express3 and Crypto Express2 is designed to meet the FIPS 140-2 Level 4 security requirements for hardware security modules.
- **CFCC Level 16 support**
 - **Hardware decimal floating point unit** for each core on the PU.
 - **Fiber Quick Connect (FQC)**, an optional feature, is a fiber harness integrated in the System z10 BC frame for a “quick” connect to ESCON and FICON LX channels.
 - **TKE 6.0 Licensed Internal Code (LIC) support**
 - **z/VM®-mode partition (LPAR) support** to contain processor types (CPs, IFLs, zIIPs, zAAPs, and ICFs)
 - **Plan ahead memory**, an optional feature, allows you to preplan to future memory upgrades. The memory upgrades can be made nondisruptively and also concurrently.
 - **Worldwide Port Name (WWPN) prediction tool**
The WWPN prediction tool assists you in preplanning and setting up your Storage Area Networks (SANs) environment prior to the installation of your System z10 server. Therefore, you can be up and running much faster after the server is installed. This tool applies to all FICON channels defined as CHPID type FCP (for communication with SCSI devices). The WWPN prediction tool is located on Resource Link.
 - **EAL5 certification**
The z10 BC server received the Common Criteria Evaluation Assurance Level 5 (EAL5) certification level (Evaluation Assurance Level 5) for the security of its LPARs that run under the control of the Processor Resource/Systems Manager™ (PR/SM™).
 - **Enhanced security using digital signatures**
Digitally Signed Firmware (Licensed Internal Code) support provided by the HMC and the SE. This support provides the following benefits:
 - Ensures that no malware can be installed on System z products during firmware updates (such as, transmission of MCL files, delivery of code loads, and restoration of critical data)
 - Designed to comply to FIPS (Federal Information Processing Standard) 140-2 Level 1 for Cryptographic LIC (Licensed Internal Code) changes.

z10 BC model

z10 BC is a single model, E10. The model naming is representative of the maximum number of customer configurable processor units (PUs) in the system. PUs are delivered in single engine increments orderable by feature code.

The following table lists some of the characteristics of the z10 BC model E10.

Table 2. System z10 BC model structure

Model	Processor Units (PUs)	Memory	I/O drawers	Maximum Channels
E10	1 to 10	4 to 248 GB	1 to 4	480

The CP features offered have varying levels of capacity. The capacity setting is based on the quantity and type of CP feature. It is identified by a **model capacity indicator**. The model capacity indicator identifies the number of active CPs rather than the total physical PUs purchased and identifies the type of capacity. The model capacity indicators are identified as A0x - Z0x, where A - Z identifies the subcapacity level and x is the number of active CP features (1 - 5).

Performance

With the expanded capacity of the z10 BC and enhancements to the I/O infrastructure, IBM continues to facilitate the consolidation of multiple servers into one z10 BC with a substantial increase in:

- Available memory
- LPARs
- Speed using InfiniBand
- Available processors in a single footprint
- 3.5 GHz high frequency z10 quad core processor chip.

IBM's Large Systems Performance Reference (LSPR) method provides comprehensive z/Architecture processor capacity ratios for different configurations of Central Processor Units across a wide variety of system control program and workload environments. For z10 BC, z/Architecture processor subcapacity indicator is defined with a A0x - Z0x notation, where x is the number of installed CPs (from one to five). There are a total of 26 subcapacity levels, designated by the letters A through Z.

For more information on LSPR, refer to <http://www.ibm.com/servers/eserver/zseries/lSpr/>.

Resource Link

Resource Link is a key component in getting your z10 BC server up and running and maintained. Resource Link provides: customized planning aids, a CHPID Mapping Tool, Customer Initiated Upgrades (CIU), power estimation tool, and education courses. Refer to Appendix B, "Resource Link," on page 135 for detailed information about Resource Link and all the functions that it can assist you with your z10 BC processor.

Fiber optic cabling

IBM Site and Facilities Services has a comprehensive set of scalable solutions to address IBM cabling requirements from product-level to enterprise level. IBM Site and Facilities Services is designed to deliver convenient, packaged services to help reduce the complexity of planning, ordering, and installing fiber optic cables. Refer to Chapter 8, "Cabling," on page 99 for additional information.

z/Architecture

The System z10 BC, like its predecessors, support 24, 31 and 64-bit addressing, as well as multiple arithmetic formats. High-performance logical partitioning via Processor Resource/System Manager (PR/SM) is achieved by industry-leading virtualization support provided by z/VM. The z/Architecture also provides key technology features such as HiperSockets and the Intelligent Resource Director, which result in a high speed internal network and an intelligent management with dynamic workload prioritization and physical resource balancing.

IBM's z/Architecture includes:

- New high-frequency z10 quad core processor chip (3.5 Ghz operation in system)
- Hardware accelerators on the chip for data compression, cryptographic functions and decimal floating point
- Integrated SMP communications
- Instructions added to z10 chip to improve compiled code efficiency
- Enablement for software/hardware cache optimization
- System z10 BC support for 1MB page frames
- Full hardware support for Hardware Decimal Floating-point Unit (HDFU)
- 64-bit general registers
- 64-bit integer instructions. Most ESA/390 architecture instructions with 32-bit operands have new 64-bit and 32- to 64-bit analogs
- 64-bit addressing is supported for both operands and instructions for both real addressing and virtual addressing
- 64-bit address generation. z/Architecture provides 64-bit virtual addressing in an address space, and 64-bit real addressing.
- 64-bit control registers. z/Architecture control registers can specify regions and segments, or can force virtual addresses to be treated as real addresses
- The prefix area is expanded from 4K to 8K bytes
- Quad-word storage consistency
- The 64-bit I/O architecture allows CCW indirect data addressing to designate data addresses above 2GB for both format-0 and format-1 CCWs
- The 64-bit SIE architecture allows a z/Architecture server to support both ESA/390 (31-bit) and z/Architecture (64-bit) guests and Zone Relocation is expanded to 64-bit for LPAR and z/VM
- 64-bit operands and general registers are used for all cryptographic instructions
- The implementation of 64-bit z/Architecture can help reduce problems associated with lack of addressable memory by making the addressing capability virtually unlimited (16 Exabytes).

z/Architecture also includes the following functions:

- ASN-and-LX-Reuse Facility
- Compare-and-Swap-and-Store Facility
- Compare-and-Swap-and-Store Facility 2
- Conditional-SSKE Facility
- Configuration-Topology Facility
- DAT-Enhancement Facility 1
- DAT-Enhancement Facility 2
- Decimal-Floating-Point Facility
- Decimal-Floating-Point-Rounding Facility
- Enhanced-DAT Facility
- ETF2-Enhancement Facility
- ETF3-Enhancement Facility
- Execute-Extensions Facility
- Extended-Immediate Facility
- Extended-I/O-Measurement-Block Facility

- Extended-I/O-Measurement-Word Facility
- Extended-Translation Facility 2
- Extended-Translation Facility 3
- Extract-CPU-Time Facility
- Fibre-channel-extensions (FCX) facility
- Floating-Point-Support-Sign-Handling Facility
- FPR-GR-Transfer Facility
- General-Instructions-Extension Facility
- HFP Multiply-and-Add/Subtract Facility
- HFP-Unnormalized-Extensions Facility
- IEEE-Exception-Simulation Facility
- List-Directed Initial Program Load
- Long-Displacement Facility
- Message-Security Assist
- Message-Security-Assist Extension 1
- Message-Security-Assist Extension 2
- Message-Security-Assist Extension 3
- Modified CCW Indirect Data Addressing Facility
- Move-With-Optional-Specifications Facility
- Multiple-Subchannel-Set Facility
- Parsing-Enhancement Facility
- PER-3 Facility
- PFPO Facility
- Restore-Subchannel Facility
- Server-Time-Protocol Facility
- Store-Clock-Fast Facility
- Store-Facility-List-Extended Facility
- TOD-Clock-Steering Facility

For more details, refer to the *z/Architecture Principles of Operation*.

Upgrade progression

zSeries z890 (Model A04) and z9 BC servers can be upgraded to a z10 BC. An upgrade includes all frames, cages, support cards, and new I/O features.

System z10 BC can be upgraded to a z10 EC Model E12.

Unsupported features/functions

This section lists the features/functions that are **not** supported on z10 BC and a recommended alternative, if applicable.

ICB-3 Links

ICB-3 links are not supported on the z10 BC.

OSA-Express features

The OSA-Express features are not supported on the z10 BC.

OSA-Express Token Ring

The OSA-Express Token Ring feature is not supported on the z10 BC.

Token Ring on the HMC

Token Ring is not available as a feature on the z10 BC Hardware Management Console (HMC).

Token Ring on the SE and TKE Workstation

Token Ring is not available as a feature on the z10 BC Support Element (SE) or Trusted Key Entry (TKE) workstation.

ICB-2 Links

ICB-2 links are not supported on the z10 BC.

ISC-3 Links in Compatibility Mode

ISC-3 compatibility mode is not supported on z10 BC.

Chapter 2. Hardware characteristics

This chapter describes the hardware features and functions for System z10 BC model E10 (machine type 2098).

Note: With the optional feature, the z10 BC server can operate on a non-raised floor.

You can also refer to the *System z10 Business Class Installation Manual for Physical Planning*, available on Resource Link at <http://www.ibm.com/servers/resourceLink>, for initial system physical planning requirements.

System frame configuration

The z10 BC frame ("A" frame) is built to Electronic Industry Association (EIA) standards. The "A" frame is shown in the following figure.

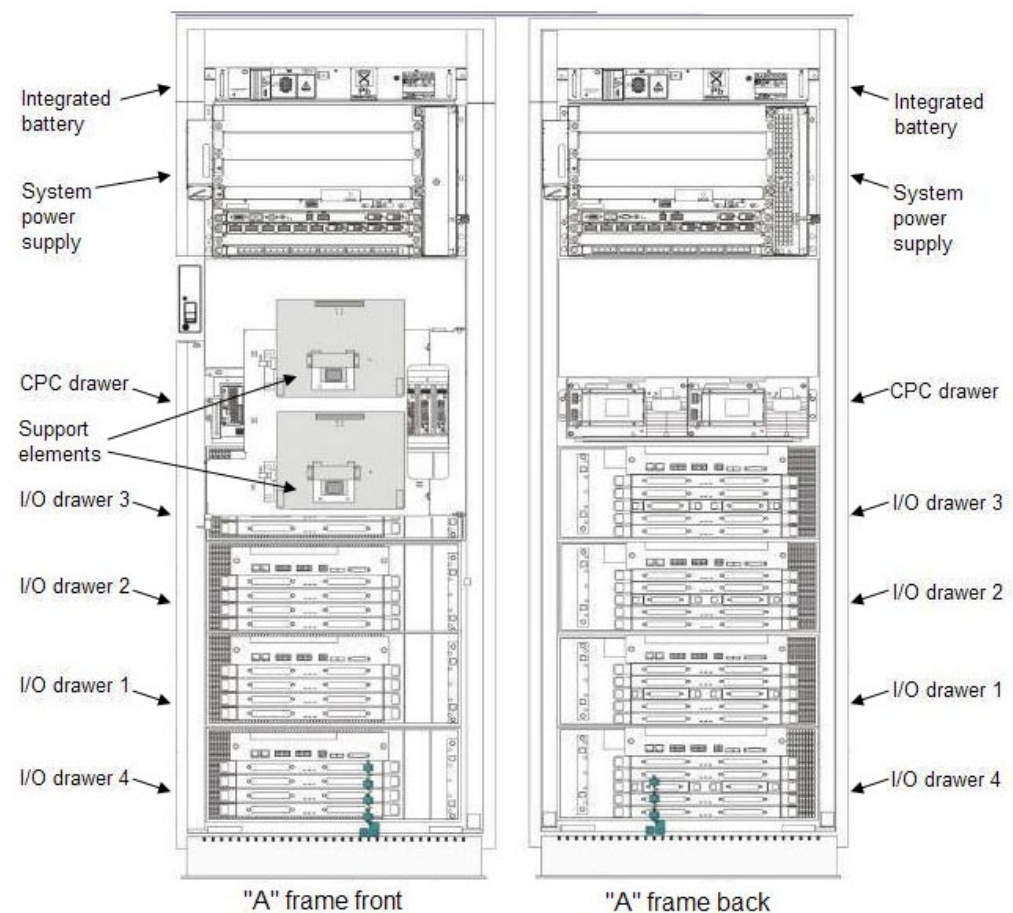


Figure 2. z10 BC frame configuration

The "A" frame consists of:

- 1 Central Processor Complex (CPC) drawer
- 0-4 Input/Output (I/O) drawer for channel attachment capability (32 I/O slots)
- Two internal Support Elements (SEs)
- System power supply

- Internal Battery Feature (IBF) - A pair of batteries installed in the top of the “A” frame for emergency backup power. Refer to the “Internal Battery Feature (IBF)” on page 21

Central Processor Complex (CPC) drawer

The Central Processor Complex (CPC) drawer is the fifth drawer in the z10 BC “A” frame. The CPC drawer consists of:

- Six SCMs (4 CP SCMs, 2 SC SCMs)
- 32 memory dimms (single or double high)
- Up to six fanout cards (HCA, MBA)
- Two OSC/ETR combination cards
- Two FSP cards
- Three DCA cards. The DCA cards are plugged directly to the CP/memory card's power board.

Figure 3 displays the contents of the CPC drawer of the z10 BC.




Figure 3. z10 BC CPC drawer

Single chip module (SCM)

The z10 BC model E10 uses a 12 PU SCM.

A PU is the generic term for the z/Architecture processor on a single chip module (SCM) that can be characterized as a:

- Central Processor (CP) to be used by the operating system
- Internal Coupling Facility (ICF) to be used by the CFCC
- Integrated Facility for Linux (IFL)
- Additional System Assist Processors (SAPs) to be used by the CSS
- IBM System z Integrated Information Processors (zIIPs)
- IBM System z Application Assist Processor (zAAP).

The z10 BC models contain the number of physical processor units (PUs) listed in Table 3 on page 13. For z10 BC you must select at least one CP, IFL, or ICF as

shown in Table 3. Any remaining PUs may be assigned as additional SAPs or additional spares, or may be assigned to optional functions such as ICFs, IFLs, CPs, On/Off CoD, or CBU engines, as indicated in Table 3. The total number of CPs, SAPs, spares, ICFs, IFLs, and CBU engines activated may not exceed the number of PUs listed for that model.

Table 3. PUs per z10 BC model

Model	PUs	Active PUs			zAAPs ²	zIIPs ²	SAPs Std	SAPs Opt	Spare PUs	Memory (GB)
		CPs ¹	ICFs ¹	IFLs ¹						
E10	10	0 - 5	0 - 10	0 - 10	0 - 5	0 - 5	2	0 - 2	0	4 to 248

Notes:

1. Only one PU (CP, ICF, or IFL) is required for any model. The total number of CPs purchased may not exceed the total number available for that model.
2. One CP must be installed with or prior to any zIIPs or zAAPs that are installed. You can purchase one zAAP and/or one zIIP for each CP on the system. This means that for every one CP, you can have one zAAP **and** one zIIP.
3. An additional 8 GB is delivered and reserved for HSA.
4. PU selection is completed by identifying the number of features when ordering.

Central Processor (CP): A Central Processor (CP) is a PU that has the z/Architecture and ESA/390 instruction sets. It can run z/VM, z/OS, z/VSE®, TPF, z/TPF, and Linux on System z operating systems and the Coupling Facility Control Code (CFCC). z10 BC processors operate only in LPAR mode; consequently all CPs are dedicated to a partition or shared between partitions. Reserved CPs can also be defined to a logical partition, to allow for nondisruptive image upgrades.

All CPs within a configuration are grouped into a CP pool. Any z/VM, z/OS, z/VSE, TPF, z/TPF, and Linux on System z operating systems can run on CPs that were assigned from the CP pool. Within the capacity of the z10 BC drawer, CPs can be concurrently added to an existing configuration permanently by using CIU or CUod, or temporarily by using On/Off CoD, CBU, and CPE.

Internal Coupling Facility (ICF): An ICF provides additional processing capability exclusively for the execution of the Coupling Facility Control Code (CFCC) in a coupling facility LPAR. Depending on the model, optional ICF may be ordered. ICFs can only be used in coupling facility logical partitions, but can be shared or dedicated, because only CFCC runs on these PUs. Software Licensing charges are not affected by the addition of ICFs. For more information, refer to “Coupling facility” on page 80.

Integrated Facility for Linux (IFL): An IFL feature provides additional processing capacity exclusively for Linux on System z workloads. An IFL can only be used in Linux on System z or z/VM LPARs. However, it can be shared or dedicated because only Linux on System z software runs on these CPs.

IFL is an optional feature for z10 BC. Up to 10 IFL features may be ordered for System z10 BC.

Software licensing charges are not affected by the addition of IFLs. For more information on software licensing, contact your IBM representative.

The IFL enables you to

- Add processing capacity dedicated to running Linux on System z on a z10 BC server.
- Run multiple Linux on System z images independently of the traditional z/Architecture, with associated savings of IBM z/Architecture.
- Define many virtual Linux on System z images on fewer real z10 BC resources.

As with any change in the LPAR configuration of a processor, the introduction of additional resources to manage may have an impact on the capacity of the existing LPARs and workloads running on the server. The size of the impact is dependent on the quantity of added resources and the type of applications being introduced. Also, one should carefully evaluate the value of sharing resources (like CHPIDs and devices) across LPARs to assure the desired balance of performance, security, and isolation has been achieved.

System z10 Applications Assist Processor (zAAP): The System z10 Application Assist Processor is a specialized processor unit that provides a Java execution environment for a z/OS environment. zAAPs are designed to operate asynchronously with the CPs to execute Java programming under control of the IBM Java Virtual Machine (JVM).

When configured with CPs within logical partitions (LPARs) running z/OS, zAAPs may help increase CP productivity and may contribute to lowering the overall cost of computing for z/OS and z/OS.e Java technology-based applications. zAAPs are designed to operate asynchronously with the CPs to execute Java programming under control of the IBM Java Virtual Machine (JVM). This can help reduce the demands and capacity requirements on CPs which may then be available for reallocation to other System z10 workloads.

The IBM JVM processing cycles can be executed on the configured zAAPs with no anticipated modifications to the Java application. Execution of the JVM processing cycles on a zAAP is a function of the Software Developer's Kit (SDK) 1.4.1 for System z10, System z9, zSeries®, z/OS, and Processor Resource/Systems Manager (PR/SM).

Note: The zAAP is a specific example of an assist processor that is known generically as an Integrated Facility for Applications (IFA). The generic term IFA often appears in panels, messages, and other online information relating to the zAAP.

z/VM 5.3 or later supports zAAPs for guest exploitation.

System z10 Integrated Information Processor (zIIP): The IBM System z10 Integrated Information Processor (zIIP) is a specialty engine designed to help improve resource optimization, enhancing the role of the server as the data hub of the enterprise. The z/OS operating system, on its own initiate or acting on the direction of the program running in SRB mode, controls the distribution of work between the general purpose processor (CP) and the zIIP. Using a zIIP can help free capacity on the general purpose processor.

z/VM 5.3 or later supports zIIPs for guest exploitation.

System Assist Processor (SAP): A SAP is a PU that runs the channel subsystem Licensed Internal Code (LIC) to control I/O operations. One of the SAPs in a configuration is assigned as a Master SAP, and is used for communication between the z10 BC drawer and the Support Element. All SAPs perform I/O operations for all logical partitions.

A standard SAP configuration provides a very well balanced system for most environments. However, there are application environments with very high I/O rates (typically some TPF environments), and in this case additional SAPs can increase the capability of the channel subsystem to perform I/O operations. Additional SAPs can be added to a configuration by either ordering optional SAPs or assigning some PUs as SAPs. Orderable SAPs may be preferred since they do not incur software charges, as might happen if PUs are assigned as SAPs.

z/VM-mode LPARs: System z10 BC allows you to define a z/VM-mode LPAR containing a mix of processor types including CPs and specialty processors (IFLs, zIIPs, zAAPs, and ICFs). This support increases flexibility and simplifies systems management by allowing z/VM 5.4 or later to manage guests to operate Linux on System z on IFLs, operate z/VSE and z/OS on CPs, offload z/OS system software overhead, such as DB2® workloads, on zIIPs, and provide an economical Java execution environment under z/OS on zAAPs, all in the same VM LPAR.

Memory

Each z10 BC CPC has its own processor memory resources. CPC processor memory can consist of both **central** and **expanded** storage.

Central storage: Central storage consists of main storage, addressable by programs, and storage not directly addressable by programs. Nonaddressable storage includes the Hardware System Area (HSA). Central storage provides:

- Data storage and retrieval for the Processor Units (PUs) and I/O
- Communication with PUs and I/O
- Communication with and control of optional expanded storage
- Error checking and correction.

Part of central storage is allocated as a fixed-sized Hardware System Area (HSA), which is not addressable by application programs. Factors affecting size are described in “Hardware System Area (HSA)” on page 16.

In z/Architecture, storage addressing is 64 bits, allowing for an addressing range up to 16 Exabytes. Consequently, all central storage in a z10 BC can be used for central storage.

Key-controlled storage protection provides both store and fetch protection. It prevents the unauthorized reading or changing of information in central storage.

Each 4 KB block of storage is protected by a 7-bit storage key. For processor-initiated store operations, access key bits 0-3 from the active program status word (PSW) are compared with bits 0-3 from the storage key associated with the pertinent 4 KB of storage to be accessed. If the keys do not match, the central processor is notified of a protection violation, the data is not stored, and a program interruption occurs. The same protection is active for fetch operations if bit 4 of the storage key (the fetch protection bit) is on. Refer to *System z10 Processor Resource/Systems Manager Planning Guide* for more information on central storage.

Expanded storage: Expanded storage can optionally be defined on System z10 servers. It is controlled by the control program, which can transfer 4 KB pages between expanded storage and central storage. The control program can use expanded storage to reduce the paging and swapping load to channel-attached paging devices in a storage-constrained environment and a heavy-paging environment.

z10 BC offers a flexible storage configuration which streamlines the planning effort by providing a single storage pool layout at IML time. The storage is placed into a single pool which can be dynamically converted to ES and back to CS as needed. Logical partitions are still specified to have CS and optional ES as before. Activation of logical partitions as well as dynamic storage reconfigurations will cause LPAR to convert the storage to the type needed.

The control program initiates the movement of data between main storage (the addressable part of central storage) and expanded storage. No data can be transferred to expanded storage without passing through main storage. With z10 BC, a **dedicated move page engine** assists in efficiently transferring data between main and expanded storage. Refer to *System z10 Processor Resource/Systems Manager Planning Guide* for more information on expanded storage.

Memory cards: Up to 32 memory cards (DIMMs) reside within a drawer. The physical card capacity can be either 2 GB (FC 1642), 4 GB (FC 1644), or 8 GB (FC 1643). Each feature code includes four DIMMs.

Note: The sum of enabled memory on each card is the amount available for use in the system.

The following list contains some general rules for memory.

- Memory cards are Field Replaceable Units (FRUs), separate from the drawer.
- Larger capacity cards may be used for repair actions and manufacturing substitution. LICCC will dial down to ordered size.
- Memory downgrades are not supported.
- Minimum memory orderable is 8 GB. Maximum memory orderable is 248 GB (+ 8 GB for HSA) when 8 GB dimms are plugged on the 32 dimm slots.
- Memory is only upgradeable in 8 GB increments between the defined minimum and maximum.
- LICCC dialing is used to offer concurrent memory upgrades within the physical memory card installed.
- The memory LICCC record for the drawer is combined with the PU LICCC record for the drawer. Both memory and PU LICCC are shipped on a single CD.

Hardware System Area (HSA): The HSA contains the CPC Licensed Internal Code (LIC) and configuration dependent control blocks. HSA is not available for program use. The HSA has a fixed size of 8 GB. Customer storage will no longer be reduced due to HSA size increase on a GA upgrade because an additional 8 GB is always delivered and reserved for HSA.

Error Checking and Correction (ECC): Data paths between central storage and expanded storage (if configured), and between central storage and the central processors and channels are checked using either parity or Error Checking and Correction (ECC). Parity bits are included in each command or data word. ECC bits are stored with data in central storage. ECC codes apply to data stored in and fetched from central storage. Memory ECC detects and corrects single bit errors. Also, because of the memory structure design, errors due to a single memory chip failure are corrected. Unrecoverable errors are flagged for follow-on action. ECC on z10 BC is performed on the memory data bus as well as memory cards.

Fanout cards

A z10 BC has one CPC drawer. The drawer includes six fanout slots. There are four main fanout cards that will plug into the z10 BC - an MBA fanout card, an HCA2-O fanout card, and HCA2-O LR fanout card, and an HCA2-C fanout card.

Memory Bus Adapter (MBA) fanouts: An MBA fanout card is used for ICB-4 attachment and provides two Self-Timed Interconnect (STIs) operating at a link rate of 2.0 GBps. The MBA fanout card is primarily used to connect a z10 BC or z10 EC to a z9 EC or z9 BC server. (However, it can connect two System z10 BC servers.)

Host Channel Adapter (HCA) fanouts: There are four types of HCA fanout cards: HCA2-C, HCA2-O, HCA2-O LR, and HCA1-O. The HCA2-C (copper) card is used for internal connection to the I/O drawer only. The HCA2-O, HCA2-O LR, and HCA1-O (optical) fanout cards are used for coupling using an InfiniBand connection. The HCA2-O fanout is designed to support a two-port 12x IB-DDR optical link operating at a link rate of 6 GBps. The HCA2-O LR fanout is designed to support a two-port 1x IB-SDR or 1x IB-DDR optical link. The 1x IB-SDR link operates at a link rate of 5.0 Gigabits per second (Gbps); the 1x IB-DDR link operates at a link rate of 2.5 Gigabits per second (Gbps). The HCA1-O fanout is designed to support a two-port 12x IB-DDR optical link operating at a link rate of 3 GBps. (The HCA1-O fanout is for z9 only.)

OSC/ETR card

Two OSC/ETR cards are required on the z10 BC. Each card provides oscillator function and External Time Reference (ETR) function.

Oscillator function: The OSC/ETR cards provide the main oscillator for z10 BC, including the STI, reference, and MBA clocks.

External Time Reference (ETR) function: Each OSC/ETR card provides attachment to the Sysplex Timer Model 2. Each OSA/ETR card connects to a different Sysplex Timer in an expanded availability configuration. Each card has an ETR port supporting an MT-RJ fiber optic connector to provide the capability to attach to a Sysplex Timer and a PPS input port supporting a coaxial connector to provide the capability to attach to the PPS output of an NTP server.

Note: The System z10 BC server ETR function does not support a multimode fiber optic cable terminated with an ESCON duplex connector. However, 62.5 micron multimode ESCON duplex jumper cables can be reused to connect to the ETR feature. This is done by installing an MT-RJ/ESCON Conversion kit between the ETR feature MT-RJ port and the ESCON duplex jumper cable.

The Sysplex Timer is a centralized time source that sets the Time-of-Day (TOD) clocks in all attached servers to maintain synchronization. The Sysplex Timer can provide the stepping signal that ensures that all TOD clocks in a multisystem environment increment in unison to permit full read/write data sharing with integrity.

If you are not using Server Time Protocol (STP), the Sysplex Timer is a mandatory hardware requirement for a sysplex consisting of more than one CPC. Additionally, the CPCs must be connected to the same External Time Reference (ETR) network.

If you are using STP, the Sysplex Timer is not required for an STP-only Coordinated Timing Network (CTN). STP is designed for servers that have been configured to be in a Parallel Sysplex or a basic sysplex (without a Coupling Facility) as well as servers that are not in a sysplex, but need to be time synchronized.

For more information on the parallel sysplex environment, refer to Chapter 6, "Sysplex functions," on page 75.

The z10 BC server, implemented in the server's Support Element Licensed Internal Code (LIC), requires the ETR Network ID of the attached Sysplex Timer Network to

be manually set in the Support Element at installation time. This function checks that the ETR Network ID being received in the timing signals via each of the server's two ETR ports matches the ETR Network ID manually set in the server's Support Element. This function provides greater checking; therefore, it helps eliminate cabling errors where either server ETR port may be incorrectly connected to Sysplex Timer units of an incorrect Sysplex Timer ETR Network. If the ETR Network ID received on one or both server ETR ports does not match the value set in the server's Support Element, that ETR port state is made semi-operational by the server. Timing signals are still received from the semi-operational ETR port, but are not used for stepping the server TOD clock. This has some important operational considerations at sysplex system IPL time as well as for running sysplex systems. Another important value of this function is that it allows verification of cabling connectivity from the Sysplex Timer to the z10 BC server prior to IPLing z/OS.

If the z10 BC server is using STP, configured in a Mixed CTN, and directly connected to the Sysplex Timer, connections to the ETR ports are required. If the z10 BC server is using STP and configured in an STP-only CTN using NTP with pulse per second as the external time source, cables are required from the PPS output of the NTP server to the PPS port on the External Time Reference (ETR) card.

Distributed Converter Assembly (DCA) cards

The Distributed Converter Assembly (DCA) cards are DC-to-DC converter cards in the CPC drawer that convert –350 volts DC to logic voltages. There are three DCA cards in the CPC drawer.

I/O drawer

The z10 BC allows up to four I/O drawers. Each I/O drawer supports ESCON, FICON, OSA Express2, OSA Express3, Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P features. An I/O drawer allows you to add channels up to the amount supported by the I/O drawer and the CPC.

Figure 4 on page 35 is an example of the I/O drawer layout for z10 BC.

I/O features

The I/O cards that are supported in the I/O drawer are shown in Table 4 on page 19. There are a total of 32 slots in the I/O drawer. You can also refer to Chapter 5, "I/O connectivity," on page 45 for more detailed information on the I/O channels and adapters.

Notes:

1. Crypto Express3, Crypto Express3-1P, Crypto Express2, and Crypto Express2-1P features use I/O slots. The Crypto Express3 feature has two PCIe adapters, the Crypto Express3-1P feature has one PCIe adapter, the Crypto Express2 feature has two PCI-X adapters, and the Crypto Express2-1P feature has one PCI-X adapter. The Crypto Express3, Crypto Express3-1P, Crypto Express2, and Crypto Express2-1P features do not have ports and do not use fiber optic cables. They are not defined in the IOCDs, and, therefore, do not receive CHPID numbers. However, they are assigned a cryptographic number.
2. ICB-4 connects directly from the STI port on an MBA fanout card on a z10 BC to an STI port on an MBA fanout card on a z9 EC or z9 BC server.
3. HCA2-O and HCA2-O LR are not I/O features. They are a fanout feature in the CPC drawer. Each fanout feature has two ports support InfiniBand coupling links.

Table 4. Channels, links, ports, and adapters summary per system

Feature	Max features	Maximum connections	Channels/ Links/ Adapters per feature	Purchase increment
16-port ESCON (FC 2323) ¹	32	480 channels	16 channels ²	4 channels
FICON Express8 10KM LX (FC 3325) ^{1, 3} FICON Express8 SX (FC 3326) ^{1, 3}	32	128 channels	4 channels	4 channels
FICON Express4 10KM LX (FC 3321) ^{1, 3} FICON Express4 4KM LX (FC 3324) ^{1, 3} FICON Express4 SX (FC 3322) ^{1, 3}	32	128 channels	4 channels	4 channels
FICON Express4-2C 4KM LX (FC 3323) ^{1, 3}	32	64 channels	2 channels	2 channels
FICON Express4-2C SX (FC 3318) ^{1, 3}	32	64 channels	2 channels	2 channels
FICON Express2 LX (FC 3319) ^{1, 3, 7} FICON Express2 SX (FC 3320) ^{1, 3, 7}	28	112 channels	4 channels	4 channels
FICON Express LX (FC 2319) ^{1, 3, 7} FICON Express SX (FC 2320) ^{1, 3, 7}	20	40 channels	2 channels	2 channels
OSA-Express3 GbE LX (FC 3362) OSA-Express3 GbE SX (FC 3363)	24	96 ports	4 ports	4 ports
OSA-Express3 10 GbE LR (FC 3370) OSA-Express3 10 GbE SR (FC 3371)	24	48 ports	2 ports	2 ports
OSA-Express3-2P GbE SX (FC 3373)	24	48 ports	2 ports	2 ports
OSA-Express3 1000BASE-T Ethernet (FC 3367)	24	96 ports	4 ports	4 ports
OSA-Express3-2P 1000BASE-T Ethernet (FC 3369)	24	48 ports	2 ports	2 ports
OSA Express2 GbE LX (FC 3364) OSA Express2 GbE SX (FC 3365) OSA Express2 1000BASE-T Ethernet (FC 3366)	24	48 ports	2 ports	2 ports
OSA-Express2 10 GbE LR (FC 3368)	24	24 ports	1 port	1 port
Crypto Express3 (FC 0864) ⁶	8	16 PCIe adptrs	2 PCIe adptrs	2 PCIe adptrs
Crypto Express3-1P (FC 0871) ⁶	8	8 PCIe adptrs	1 PCIe adptr	1 PCIe adptr
Crypto Express2 (FC 0863) ⁶	8	16 PCI-X adptrs	2 PCI-X adptrs	2 PCI-X adptrs
Crypto Express2-1P (FC 0870) ⁶	8	8 PCI-X adptrs	1 PCI-X adptr	1 PCI-X adptr
ISC-3 ¹	12	48 links ⁴	4 links	1 link
ICB-4 ^{1, 5}	6	12 links ⁴	2 links	1 link
IFB (12x IB-DDR) ^{1, 5}	6	12 links ⁴	2	2 links
IFB (1x IB-DDR) ^{1, 5}	6	12 links ⁴	2 links	2 links

Table 4. Channels, links, ports, and adapters summary per system (continued)

Feature	Max features	Maximum connections	Channels/ Links/ Adapters per feature	Purchase increment
Notes: <ol style="list-style-type: none"> 1. A minimum of one I/O feature (ESCON or FICON) or one coupling link (ICB-4, ISC-3, IFB) is required. 2. Each ESCON feature has 16 channels, of which a maximum of 15 may be activated. One is reserved as a spare. 3. A maximum number of combined FICON Express8, FICON Express4, FICON Express2, and FICON Express features is 32. 4. A maximum number of coupling links combined cannot exceed 64 per server. (ICB-4, active ISC-3, and IFB links) 5. ICB-4s and IFBs are not included in the maximum feature count for I/O slots, but they are included in the CHPID count. 6. An initial order for Crypto Express3 is four PCIe adapters (two features), an initial order for Crypto Express3-1P is two PCIe adapters (two features), an initial order for Crypto Express2 is four PCI-X adapters (two features), and an initial order for Crypto Express2-1P is two PCI-X adapters (two features). Each PCI-X adapter or PCIe adapter can be configured as either a coprocessor or an accelerator. 7. FICON Express and FICON Express2 features can be carried forward on an upgrade, but they cannot be ordered. 				

Distributed Converter Assembly (DCA) cards

The Distributed Converter Assembly (DCA) cards are DC-to-DC converter cards in the I/O drawer that convert –350 volts DC to logic voltages. There are two DCA cards in each I/O drawer.

PSC24V card

The PSC24V card is a power sequence control (PSC) card used to turn on/off specific control units from the CPC. The PSC24V card in the I/O drawer provides the physical interface between the cage controller and the PSC boxes, located outside the I/O drawer in the system frame. Only one PSC24V card is required on the z10 BC server. Each card has two jacks that are used to connect to the PSC boxes.

Note: The PSC24V card is **not** hot pluggable.

For more information on PSC, refer to “Power sequence control” on page 26.

Support Element

The z10 BC is supplied with two integrated laptop computers that function as a primary and alternate Support Elements (SEs). The SEs are ThinkPads positioned over each other in the front of the “A” frame. The Support Elements communicate with the CPC and each other through the service network. The second, or alternate SE is designed to function as a backup and to preload SE Licensed Internal Code. Power for the SE is supplied by the server frame.

The SE contains the following:

- Licensed Internal Code for the CPC.
- Hardware system definitions for the CPC (contained in the reset, image, and load profiles for the CPC and IOCDs).
- Battery-powered clock used to set the CPC time-of-day (TOD) clock at power-on reset. In ETR timing mode, CPC attachment to a Sysplex Timer causes the CPC

TOD clock to take its time setting from the Sysplex Timer. In STP timing mode, the CPC TOD clock is initialized to Coordinated Server Time (CST).

- Two 10/100 Mbps Ethernet LAN adapters to connect the SE to the Hardware Management Console through an external ethernet switch. An ethernet switch (FC 0089) is required to provide an ethernet connection between the Hardware Management Console and the internal SEs.

An ethernet switch is **NOT** shipped automatically with your order.

- An ethernet LAN adapter or LAN on board to connect the SE to the CPC through the power service network.

For more detailed information on the Support Element, refer to the Chapter 9, "Hardware Management Console and Support Element," on page 105 or to the *System z10 Support Element Operations Guide*.

Support Element attachment to HMC

The Support Elements have dual ethernet LAN adapters used to communicate with the HMC. This is offered as part of the initial order or as a Manufacturing Engineering Specification (MES). An ethernet switch (FC 0089) is required to provide an ethernet attachment between the HMC and the internal SEs. An ethernet switch is shipped automatically on every order unless you deselected FC 0089.

System power supply

The system power supply located in the top of the "A" frame provides the control structure to support the z10 BC power requirements for the CPC drawer and four I/O drawers.

The z10 BC power subsystem basic components include:

- Bulk Power Assembly (BPA) - provides the prime power conversion and high voltage DC distribution.
- Bulk Power Controller (BPC) - is the main power controller and cage controller for the BPA.

The BPC is the principal control node for the z10 BC diagnostic/service and power/cooling system. It is the cage controller for the BPA cage and connects to both ethernet service networks.

- Bulk Power Distribution (BPD) - distributes -350 VDC and RS422 communications to logic cage power Field Replaceable Units (FRUs)
- Bulk Power Fan (BPF) - is a cooling device
- Bulk Power Regulator (BPR) - is the main front end power supply that converts line voltage to regulated -350 VDC
- Bulk Power Enclosure (BPE) - is the metal enclosure that contains the back plane
- Bulk Power Hub (BPH) - provides ethernet connection to the Support Element and all cage controllers in the system.
- Distributed Converter Assemblies (DCAs).

Internal Battery Feature (IBF)

The optional Internal Battery Feature (IBF) (FC 3211) provides the function of a local uninterruptible power source. It has continuous self-testing capability for battery backup which has been fully integrated into the diagnostics, including Remote Service Facility (RSF) support.

The IBF provides battery power to preserve processor data in case of a loss of power on both of the AC supplies from the utility company. The IBF can hold power from seven to twenty minutes depending upon the system configuration that will protect system operation for the majority of utility power interruptions.

IBFs must be ordered as a pair. The pair equals one front and one back battery. The IBF is fully integrated into the server power control/diagnostic system that provides full battery charge, and test and repair diagnostics. For more information on the IBF, refer to *System z10 Business Class Installation Manual for Physical Planning*.

Internet Protocol Version 6

IPv6 is the protocol designed by the Internet Engineering Task Force (IETF) to replace Internet Protocol Version 4 (IPv4) to satisfy the demand for additional IP addresses. IPv6 expands the IP address space from 32 bits to 128 bits enabling a far greater number of unique IP addresses.

The HMC and supports element are designed to support customer internal and open networks that are configured to use only IPv6 addresses, only IPv4 addresses, or a combination of the two.

Multiple Subchannel Sets (MSS)

Multiple Subchannel Sets (MSS) provides a second set of subchannels for defining Parallel Access Volume (PAV) aliases. This function provides relief from the 64K device limit by allowing PAV aliases to be defined without making device numbers unavailable for defining additional devices. This is designed to provide an I/O constraint benefit. Two subchannel sets per Logical Channel Subsystem (LCSS) are designed to enable a total of 63.75K subchannels in set-0 and the addition of 64K-1 subchannels in set-1. MSS is designed to provide greater I/O device configuration capabilities for larger enterprises. MSS is supported by ESCON (CHPID type CNC), FICON (when configured as CHPID type FC or FCV), and z/OS.

LPAR mode

LPAR mode is the mode of operation for the z10 BC. It allows you to:

- Define ESA/390, ESA/390 TPF, coupling facility, z/VM-mode, and Linux-only logical partitions
- Define and use up to 248 GB in a single LPAR.
- Dynamically reconfigure storage between logical partitions.

You can define and activate up to 30 logical partitions for each CPC.

After you define and activate an ESA/390 or ESA/390 TPF logical partition, you can load a supporting operating system into that logical partition.

Processor Resource/Systems Manager (PR/SM) enables logical partitioning of the CPC.

Resources for each logical partition include:

- Processor units (CPs, ICFs, IFLs, zIIPs, or zAAPs)
- Storage (central storage and expanded storage)
- Channels.

Processor units

On z10 BC, PUs can be used within a logical partition as Central Processors (CPs), Internal Coupling Facilities (ICFs), Integrated Facilities for Linux (IFLs), System z10 Integrated Information Processor (zIIP), or System z10 Application Assist Processors (zAAPs). The initial allocation of CPs, ICFs, IFLs, zIIPs, and zAAPs to a logical partition is made when the logical partition is activated.

Within a logical partition on z10 BC, they may be used as follows:

- CPs can be dedicated to a single logical partition or shared among multiple logical partitions. The use of CP resources shared between logical partitions can be limited and modified by operator commands while the logical partitions are active. CPs that are dedicated to a logical partition are available only to that logical partition.
- ICFs, IFLs, zIIPs, and zAAPs are available as orderable features on z10 BC for use in a logical partition. ICFs are available as a feature for use in a coupling facility (CF) logical partition (refer to “Internal Coupling Facility (ICF)” on page 13 for additional information). IFLs are available as a feature for running Linux on System z. zAAPs are available as a feature for providing special purpose assists that execute JAVA programming under control of the IBM Java Virtual Machine (JVM) (refer to “System z10 Applications Assist Processor (zAAP)” on page 14 for additional information).

Storage

Before you can activate logical partitions, you must define central storage and optional expanded storage to the logical partitions. Refer to “Central storage” on page 15 and “Expanded storage” on page 15 for more information.

All installed storage is initially configured as central storage. Each individual logical partition is limited to a maximum of 2 GB (in ESA/390 mode) or 256 GB (in z/Architecture mode) of central storage. When a logical partition is activated, the storage resources are allocated in contiguous blocks.

For z10 BC, logical partition central storage granularity is a minimum of 128 MB and increases as the amount of storage defined for the logical partition increases. You can dynamically reallocate storage resources for z/Architecture and ESA/390 architecture logical partitions using **Dynamic Storage Reconfiguration**. Dynamic storage reconfiguration allows both central and expanded storage allocated to a logical partition to be changed while the logical partition is active. It provides the capability to reassign storage from one logical partition to another without the need to POR the CPC or IPL the recipient logical partition. For more information, refer to *System z10 Processor Resource/Systems Manager Planning Guide*.

Note: You cannot share allocated central storage or expanded storage among multiple logical partitions.

Expanded storage granularity for logical partitions is fixed at 128 MB.

Channels

You can allocate channels to logical partitions as follows:

- **Dedicated channels**

Dedicated channels are unshared channels and can only be used by one logical partition. All channel types supported by the model can be allocated as dedicated channels.

- **Reconfigurable channels**

Reconfigurable channels are unshared channels that can be moved among logical partitions within an LCSS but can only belong to one logical partition at a given time. All channel types supported by the model can be allocated as reconfigurable channels.

- **Shared channels**

The Multiple Image Facility (MIF) allows channels to be shared among multiple logical partitions in a Logical Channel Subsystem (LCSS). Shared channels are configured to a logical partition giving the logical partition a channel image of the shared channel that it can use. Each channel image allows a logical partition to independently access and control the shared channel as if it were a physical channel assigned to the logical partition. For more information, refer to “Multiple Image Facility (MIF)” on page 40.

You can define the channels, shown in Table 7 on page 32, as shared among multiple logical partitions within an LCSS so that the shared channels can be accessed by more than one logical partition in an LCSS at the same time.

On z10 BC with coupling facility logical partitions, CFP, CBP, and ICP channels may be shared by many ESA logical partitions and one coupling facility logical partition.

- **Spanned channels**

Spanned channels are channels that are configured to multiple Logical Channel Subsystems (LCSSs) and are transparently shared by any or all of the configured LPARs without regard to the LCSS to which the LPAR is configured.

- **Device Sharing**

You can share a device among logical partitions by:

- Using a separate channel for each logical partition
- Using a shared channel
- Using a spanned channel.

LPAR time offset support

Logical partition time offset support provides for the optional specification of a fixed time offset (specified in days, hours, and quarter hours) for each logical partition activation profile. The offset, if specified, will be applied to the time that a logical partition will receive from connected Sysplex Timers or from the Current Time Server (CTS) in a Coordinated Timing Network (CTN).

This support can be used to address the customer environment that includes multiple local time zones with a single Sysplex Timer, two Sysplex Timers in an expanded availability configuration, or a Current Time Server (CTS) in a CTN.

It is sometimes necessary to run multiple Parallel Sysplexes with different local times and run with the time set to GMT=LOCAL. This causes the results returned in the store clock (STCK) instruction to reflect local time. With logical partition time offset support, logical partitions on each z10 BC CPC in a Parallel Sysplex that need to do this can specify an identical time offset that will shift time in the logical partition sysplex members to the desired local time. Remaining logical partitions on the z10 BC CPCs can continue to participate in current date production Parallel Sysplexes utilizing the same Sysplex Timer(s) or CTS with the time provided by the Sysplex Timer(s) or CTS.

This function is supported by all in service releases of z/OS.

For more information on logical partitions, refer to the *System z10 Processor Resource/Systems Manager Planning Guide* and to the *System z Input/Output Configuration Program User's Guide for ICP IOCP*.

Server Time Protocol (STP)

Server Time Protocol (STP) (FC 1021) provides the means for multiple System z10, System z9, z890, and z990 servers to maintain time synchronization with each other without using a Sysplex Timer. STP is designed to synchronize servers configured in a Parallel Sysplex or a sysplex without a coupling facility, as well as servers that are not in a sysplex.

STP uses a message-based protocol to transmit timekeeping information over externally defined Coupling Links between servers. Unlike the Sysplex Timer, which distributes time to multiple servers in a star pattern, STP distributes time messages in layers (called stratus). The timekeeping information is needed to determine the Coordinated Server Time (CST) at each server. The Coupling Links used to transport STP messages include ISC-3 links configured in peer mode, ICB-4 links, and IFB links. These links can be the same links already being used in a Parallel Sysplex for coupling facility communications.

For more details about Server Time Protocol, refer to “Server Time Protocol (STP)” on page 83.

For hardware and software requirements, refer to the STP website located at <http://www.ibm.com/systems/z/pos/stp.html>.

Hardware Management Console (HMC)

On z10 BC, the Hardware Management Console (HMC) is a desktop PC that controls and monitors status for the Central Processor Complexes (CPCs). The HMC provides a single point of control and single system image for those CPCs defined to it. One Hardware Management Console can control up to 100 CPCs. One CPC can be controlled by up to 32 Hardware Management Consoles.

The HMC supports both single and dual ethernet configurations. With dual ethernet, the HMC is used to control and monitor multiple systems, local or remote, on an Ethernet local area network (LAN) using the Hardware Management Console Application (HWMCA). The HMC is supplied with two Ethernet ports.

The physical location of the Hardware Management Console hardware features (standard and/or optional) are dictated by the specific PC. Some features may be mutually exclusive with other features depending on the PC model. Each CPC must be connected to at least one Hardware Management Console on the same network as the SEs of the CPC.

For more detailed information on the Hardware Management Console, refer to Chapter 9, “Hardware Management Console and Support Element,” on page 105 or to the *System z Hardware Management Console Operations Guide*.

Bolt-down kit

A bolt-down kit is available for a low raised floor installation (9 to 13 inches), a high raised floor installation (12 to 22 inches), and a non-raised floor installation. You will need to order only one bolt-down kit.

Power sequence control

The optional power sequence controller (PSC) is available on the z10 BC. The PSC feature provides the ability to turn on/off specific control units from the CPC. The PSC feature consists of two PSC boxes, one PSC24V card, and with PSC cables.

You can order one PSC feature on your System z10 BC.

Additional features/functions supported

In addition to the standard and optional features previously listed, the design of the z10 BC also provides the following functions:

Monitoring and estimating CPC power consumption and temperature

You can monitor the power consumption and the internal temperature of a specific CPC using the Hardware Management Console (HMC) or the Active Energy Manager. In addition to providing the power consumption and temperature of a specific CPC, the Active Energy Manager also provides the aggregated temperature and power for a group of systems or a complete data center. Active Energy Manager can display this data in a format that shows trends over a specified time intervals. It is a plug-in to IBM Director.

Before using Active Energy Manager, you must enable the SNMP APIs and define a community name for Active Energy Manager. This is specified on the Customize API Settings task on the HMC. Once you have configured the SNMP support on the HMC, you must set up Active Energy Manager so it can communicate to the HMC. You can do this, within Active Energy Manager, by defining it as an SNMP device. Once this is complete, the Active Energy Manager can communicate to the HMC.

For more information, refer to the IBM Systems Software Information Center website (<http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp>). Expand **IBM Systems Software Information Center** located in the left-hand navigation pane, select **Product listing**, then select **IBM Director extension: Active Energy Manager** from the product listing.

You can estimate the power consumption of a specific z10 EC machine model and its associated configuration using the Power Estimation tool. The exact power consumption for your machine will vary. The purpose of the tool is to produce an estimation of the power requirements to aid you in planning for your machine installation. This tool is available on Resource Link.

Preplanning and setting up the Storage Area Network (SAN) environment

The WWPN prediction tool assists you in preplanning and setting up your Storage Area Networks (SANs) environment prior to the installation of your System z10 server. Therefore, you can be up and running much faster after the server is installed. This tool applies to all FICON channels defined as CHPID type FCP (for communication with SCSI devices). The WWPN prediction tool creates WWPN assignments that are required to set up your SAN and creates a binary configuration that can be imported by your system.

The WWPN prediction tool is located on Resource Link.

Chapter 3. Software support

This chapter describes the software support for the System z10 BC. The z10 BC supports LPAR mode only. The following table displays a summary of the minimum supported operating systems levels for the z10 BC.

Table 5. Supported operating systems for z10 BC

Operating System	ESA/390 (31-bit)	z/Architecture (64-bit)
z/OS Version 1 Releases 9, 10, 11	No	Yes
z/OS Version 1 Release 7 ¹ and 8 ² with IBM Lifecycle Extension for z/OS V1.7 and V1.8	No	Yes
Linux on System z ^{3, 8} : Red Hat RHEL 4 and Novell SUSE SLES 9	Yes	Yes
Linux on System z ^{3, 8} : Red Hat RHEL 5 and Novell SUSE SLES 10, 11	No	Yes
z/VM Version 5 Release 3 ^{4, 9} and 4 z/VM Version 6 Release 1	No ⁵	Yes
z/VSE Version 4 Release 1 ^{3, 6, 2 6} , and 3 ^{6, 7}	No	Yes
z/TPF Version 1 Release 1	No	Yes
TPF Version 4 Release 1 (ESA mode only)	Yes	No
Notes: <ol style="list-style-type: none"> 1. z/OS V1.7 support was withdrawn September 30, 2008. The Lifecycle Extension for z/OS V1.7 (5637-A01) makes fee-based corrective service for z/OS V1.7 available through September 2010. With this Lifecycle Extension, z/OS V1.7 supports the z10 Server. Certain functions and features of the z10 Server require later releases of z/OS. For a complete list of software support, see the 2098DEVICE Preventive Planning (PSP) bucket. 2. z/OS V1.8 support was withdrawn September 30, 2009. The Lifecycle Extension for z/OS V1.8 (5638-A01) makes fee-based corrective service for z/OS V1.8 available through September 2011. With this Lifecycle Extension, z/OS V1.8 supports the z10 Server. Certain functions and features of the z10 Server require later releases of z/OS. For a complete list of software support, see the 2098DEVICE Preventive Planning (PSP) bucket. 3. Compatibility support for listed releases. Compatibility support allows OS to IPL and operate on z10. 4. Requires compatibility support, which allows z/VM to IPL and operate on the z10 providing System z9 functionality for the base OS and guests. 5. z/VM supports 31-bit and 64-bit guests. 6. z/VSE V4 is designed to exploit 64-bit real memory addressing, but does not support 64-bit virtual memory addressing. 7. z/VSE V4.3 preview announcement October 20, 2009. 8. RHEL is an abbreviation for Red Hat Enterprise Linux. SLES is an abbreviation for SUSE Linux Enterprise Server. 9. z/VM 5.3 reaches end-of-service on September 30, 2010. 		

This information applies to z10 BC CPCs running in an LPAR mode.

Any program written for z/Architecture or ESA/390 architecture mode can operate on CPCs operating in the architecture mode for which the program was written, provided that the program:

- Is not time-dependent.
- Does not depend on the presence of system facilities (such as storage capacity, I/O equipment, or optional features) when the facilities are not included in the configuration.
- Does not depend on the absence of system facilities when the facilities are included in the configuration.

- Does not depend on results or functions that are defined as unpredictable or model dependent in the *z/Architecture Principles of Operation* or in the *Enterprise System Architecture/390 Principles of Operation*.
- Does not depend on results or functions that are defined in this publication (or, for logically partitioned operation, in the *System z10 Processor Resource/Systems Manager Planning Guide*) as being differences or deviations from the appropriate *Principles of Operation* publication.
- Does not depend on the contents of instruction parameter fields B and C on interception of the SIE instruction.

Any problem-state program written for ESA/390 architecture mode can operate in z/Architecture mode provided that the program complies with the limitations for operating in ESA/390 mode and is not dependent on privileged facilities which are unavailable on the system.

For z/OS to operate as a guest of z/VM on a z10 BC, z/OS and z/VM must be operating in 64-bit mode. z/VM V5 and z/VM V6 are only supported running in 64-bit mode.

z/OS

z/OS V1R8, V1R9, V1R10, and V1R11 provide support for System z10 BC. Refer to the following table for a list of the supported functions and releases.

	V1R8	V1R9	V1R10	V1R11
CPU Measurement facility	x ¹	x ¹	x	x
InfiniBand coupling	x ¹	x ¹	x ¹	x
CFCC Level 16	x ¹	x ¹	x ¹	x
Communicate QDIO Queue Counts	x ¹	x ¹	x	x
HCD support	x ¹	x ¹	x	x
OSA-Express3 features	x ²	x ²	x	x
Capacity provisioning		x	x	x
64-way single system image		x	x	x
Large page (1 MB) storage		x	x	x
FICON Express8 support	x ¹	x ¹	x ¹	x
Optimized latency mode				x ¹
Notes:				
1. Requires PTFs.				
2. z/OS V1.8 or V1.9 with PTFs is required to support CHPID type OSD and the exploitation of 4 ports per OSA-Express3 GbE LX feature, OSA-Express3 GbE SX feature, and OSA-Express3 1000BASE-T Ethernet feature; and the exploitation of 2 ports per OSA-Express3-2P GbE SX feature and OSA-Express3-2P 1000BASE-T Ethernet feature.				

Refer to the z/OS subset of the 2098DEVICE Preventive Service Planning (PSP) bucket prior to installing a z10 BC server.

z/VM

z/VM V5.3, V5.4, and V6.1 provide support for System z10. Refer to the following table for a list of the supported functions and releases.

	V5R3	V5R4	V6R1
Installing Linux on System z from the HMC		x	x
z/VM Systems Management from the HMC	x	x	x
Enhanced Systems Management from the HMC		x	x
QDIO connection isolation support	x ¹	x ¹	x
Select z/VM system and guest exploitation of System z10		x	x
z/VM system and guest exploitation of z9 functionality level	x ¹	x	x
Ability to define, modify, and delete a coupling connection using an InfiniBand link, CHPID type CIB, when z/VM is the controlling LPAR for dynamic I/O	x ¹	x	x
Support for CHPID type OSD and the exploitation of 4 ports per OSA-Express3 GbE LX feature, OSA-Express3 GbE SX feature, and OSA-Express3 1000BASE-T Ethernet feature; and the exploitation of 2 ports per OSA-Express3-2P GbE SX feature and OSA-Express3-2P 1000BASE-T Ethernet feature	x ¹	x	x
FICON Express8 support	x	x	x
Exploitation of the System z10 server cache management instructions			x
CFCC level 16	x ¹	x	x
Optimized latency mode	x ¹	x ¹	x ¹
QDIO port isolation	x ¹	x ¹	x
Notes: 1. Requires PTFs.			

Refer to the z/VM subset of the 2098DEVICE Preventive Service Planning (PSP) bucket prior to installing a z10 BC or IPLing a z/VM image.

z/VSE

z/VSE V3.1, z/VSE V4.1, and z/VSE V4.2 provide support for System z10 BC.

Compatibility support for z/VSE V3.1 and z/VSE V4.1 is provided with PTFs. Compatibility support allows z/VSE to IPL and operate on the System z10 BC, both in an LPAR and as a guest under z/VM.

z/VSE 4.1 with PTFs or z/VSE 4.2 is required to support CHPID type OSD and the exploitation of 4 ports per OSA-Express3 GbE LX feature, OSA-Express3 GbE SX feature, and OSA-Express3 1000BASE-T Ethernet feature.

z/VSE 4.1 and later support FICON Express8.

Refer to the z/VSE subset of the 2098DEVICE Preventive Service Planning (PSP) bucket prior to installing the z10 BC server.

Linux on System z

Linux on System z (Novell SUSE SLES 9, 10, 11; Red Hat RHEL 4, 5) provides compatibility support for System z10 BC.

You can run Linux on System z natively in an LPAR or as a guest under z/VM. You can also isolate your Linux on System z in its own workspace using the Integrated Facility for Linux (IFL). IFL is a hardware option for capacity dedicated to Linux on System z workload. The IFL enables you to:

- Add processing capacity dedicated to running Linux on System z on a System z10 BC server
- Run multiple Linux on System z images independently of the traditional z/Architecture, with associated savings of IBM z/Architecture
- Define many virtual Linux on System z images on fewer real System z10 BC resources.

TPF

TPF V4.1 and z/TPF V1.1 provide support for System z10 BC.

z/TPF V1.1 is required to support 64 engines per z/TPF LPAR.

TPF V4.1 and z/TPF V1.1 support CDLC for CCL V2 and the OSA-Express2 OSN card.

TPF V4.1 at PUT 13 with PTFs is required to support CHPID type OSD and the 2 ports per OSA-Express3 feature.

z/TPF V1.1 supports up to 64 engines per z/TPF LPAR, and supports CEX2A, but not CEX2C. TPF 4.1 does not support Crypto Express2.

z/TPF V1.1 at PUT 4 with APARs is required to support CHPID type OSD and the exploitation of 4 ports per OSA-Express3 GbE LX feature and OSA-Express3 GbE SX feature.

z/TPF V1.1 with PTFs supports Crypto Express3 accelerator.

Chapter 4. Channel subsystem structure

A channel subsystem (CSS) structure for z10 BC is designed for 256 channels. With the scalability benefits provided by System z10 BC, it is essential that the channel subsystem (CSS) structure is also scalable and permits “horizontal” growth. This is facilitated by allowing more than one logical channel subsystem (LCSS) on a single z10 BC.

Table 6. Channel, port, adapter maximums

Type	z10 BC Maximum
ESCON	32 cards / 480 channels
FICON Express ¹	20 features / 40 channels
FICON Express2 ¹	20 features / 80 channels
FICON Express4 ¹	32 features / 128 channels
FICON Express8 ¹	32 features / 128 channels
OSA-Express2 ²	24 features / 48 ports
OSA-Express3 ²	24 features / 96 ports
IC link	32 links
ISC-3 link ³	12 mother cards / 48 links
IFB link ³	12 links
ICB-4 link ³	12 links
Crypto Express2 ^{5, 6}	8 cards/16 adapters
Crypto Express2-1P ^{5, 6}	8 cards/8 adapters
Crypto Express3 ^{5, 6}	8 cards/16 adapters
Crypto Express3-1P ^{5, 6}	8 cards/8 adapters
Notes: <ol style="list-style-type: none"> Maximum number of combined FICON Express8, FICON Express4, FICON Express2, and FICON Express features is 32. Maximum number of combined OSA-Express2 and OSA-Express3 features is 24. Maximum number of combined ISC-3 links, IFB links, and ICB-4 links is 64. Each coupling feature cannot exceed its individual maximum limit (shown in the table). The maximum number of combined FICON, OSA, and Crypto features cannot exceed 60. The maximum number of combined Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P features is eight. The initial order for Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P is two features. 	

The CSS structure offers the following:

- Two logical channel subsystems (LCSSs)
 - Each LCSS can have up to 256 channels defined
 - Each LCSS can be configured with one to 15 logical partitions.
- Spanned channels are shared among logical partitions across LCSSs. For more information on spanned channels, refer to Table 7 on page 32 and to “Spanned channels” on page 41.

Note: One operating system image continues to support up to a maximum of 256 Channel Path Identifiers (CHPIDs).

The I/O Subsystem (IOSS) continues to be viewed as a single Input/Output Configuration Data Set (IOCDS) across the entire system with up to two LCSSs. Only one Hardware System Area (HSA) is used for the multiple LCSSs.

A CHPID is a two-digit hexadecimal number that identifies a channel path in the CPC. A Physical Channel Identifier (PCHID) is a three-digit number that identifies the physical location (drawer, slot, card port) for a channel path in the CPC. An adapter ID (AID) is a two-digit hexadecimal number that identifies HCA2-O or HCA2-O LR fanout card location. CHPIDs are associated with ports on an adapter and the AID is used in that definition.

The CHPID Mapping Tool can help you map your PCHIDs to the CHPID definitions in your IOCP source statements. The tool will provide you with a new report with your CHPID assignment in addition to the PCHID values. The CHPID Mapping Tool is available from Resource Link, <http://www.ibm.com/servers/resourceLink>, as a standalone PC-based program. For more information on the CHPID Mapping Tool, CHPIDs, PCHIDs or AIDs, refer to *System z CHPID Mapping Tool User's Guide*.

IOCP channel, link, and adapter definitions

The following table lists the channel types as defined in an IOCDS that are used with z10 BC systems.

Table 7. Channels, links, and adapters with CHPID type

Channels/Links/Adapters	CHPID type	May be defined as Shared	May be defined as Spanned
ESCON channels:			
Connection Channel (ESCON architecture)	CNC	yes	no
Channel-to-Channel (connects to CNC)	CTC	yes	no
ESCON channels connected to converter:			
Conversion Channel (ESCON to Parallel Block Multiplexer (BL))	CVC	no	no
Conversion Channel (ESCON to Parallel Byte Multiplexer (BY))	CBY	no	no
FICON Express. A FICON channel that attaches to an ESCON Directory Model 5.	FCV	yes	no
FICON channels — native FICON, zHPF, or CTC for attachment to FICON channels on System z servers, directors, control units, and printers	FC	yes	yes
Fibre Channel Protocol (FCP) for communicating with SCSI devices	FCP	yes	yes
ISC-3 peer mode links (connects to another ISC-3)	CFP	yes	yes
ICB-4 peer links (connects to another ICB-4)	CBP	yes	yes
IC peer links (connects to another IC)	ICP	yes	yes
IFB peer links (connects to another IFB)	CIB	yes	yes
HiperSocket	IQD	yes	yes
OSA adapters using QDIO architecture	OSD	yes	yes
OSA adapters using non-QDIO architecture for TCP/IP and/or SNA/APPN/HPR traffic	OSE	yes	yes

Table 7. Channels, links, and adapters with CHPID type (continued)

Channels/Links/Adapters	CHPID type	May be defined as Shared	May be defined as Spanned
OSA 1000BASE-T Ethernet adapters for TN3270E, non-SNA DFT, IPL CPCs, and LPARs, OS system console operations	OSC	yes	yes
OSA-Express for NCP: NCPs running under IBM Communication Controller for Linux (CDLC)	OSN	yes	yes

Each of these channel types requires that a CHPID be defined, even if it is an internal channel and no physical hardware (channel card) exists. Each channel, whether a “real” channel or a virtual (such as HiperSocket) must be assigned a unique CHPID within the LCSS. You can arbitrarily assign a number within the X'00' to X'FF' range. Real channels require a PCHID value to be defined. Most of these channel types can be shared and used concurrently among multiple LPARs within the same LCSS. Refer to “Multiple Image Facility (MIF)” on page 40 for more information on shared channels.

PCHIDs are used for ICB-4 connections. AIDs are used for InfiniBand connections.

Coupling link peer channels

You may define an ISC-3 feature as CFP, an ICB-4 as CBP, and an IFB link as CIB. Any available/unused CHPID may be defined as ICP.

You can configure a CFP, CBP, ICP, or CIB channel path as:

- An unshared dedicated channel path to a single logical partition.
- An unshared reconfigurable channel path that can be configured to only one logical partition at a time but which can be dynamically moved to another logical partition by channel path reconfiguration commands. Reconfigurable support for CFP, CBP, CIB, and ICP is limited to two coupling facility logical partitions total. One coupling facility logical partition in the initial access list and one other coupling facility partition in the candidate list.
- A shared or spanned channel path that can be concurrently used by the logical partitions to which it is configured. A peer channel cannot be configured to more than one coupling facility logical partition at a time, although it can be configured to multiple z/Architecture or ESA/390 logical partitions in addition to the single coupling facility logical partition.
- Timing-only links. These are coupling links that allow two servers to be synchronized using Server Time Protocol (STP) messages when a coupling facility does not exist at either end of the coupling link.

Note: ICP CHPIDs are not supported for a timing connection.

Each ICP channel path must specify which ICP channel path it is logically connected to.

The System z9 and System z10 models support dynamic I/O configuration for all peer channel path types.

Subchannel connectivity

With two Logical Channel Subsystems comes more subchannels. There is a maximum of 65280 subchannels per LCSS in subchannel set 0 and 65535 subchannels per LCSS in subchannel set 1. With two Logical Channel Subsystems, each LCSS can have its own set of subchannels in both subchannel sets and each logical partition in each LCSS can have access to the subchannels in both sets of its LCSS.

With two Logical Channel Subsystems you can have:

- Up to a maximum of 65280 devices/subchannels per LCSS for subchannel set 0
- Up to a maximum of 65535 devices/subchannels per LCSS for subchannel set 1
- Up to a maximum of 261630 devices for two LCSSs (two times the maximum devices/subchannels for subchannel set 0 and 1 ($2 * (65280 + 65535)$)).

Each LPAR can access all the devices in its assigned LCSS.

This capability relieves the I/O device configuration constraints experienced by large system configurations.

Guidelines for maximum availability

When configuring devices with multiple paths to the same CPC, select any of the channel paths from any I/O card shown in Figure 5 on page 37 that:

- Are available on the CPC you are defining
- Are the correct type (FICON, ESCON, etc.) to meet the control unit, coupling facility, or network attachment requirements
- Satisfy the rules regarding the mixing of channel types to a control unit.


Card slot 01 - DCA

Card slot 02 - Domain 0

Card slot 03 - Domain 1

Card slot 04 - Domain 1

Card slot 05 - Domain 0



I/O drawer front view

I/O drawer rear view

Card slot 06 - DCA

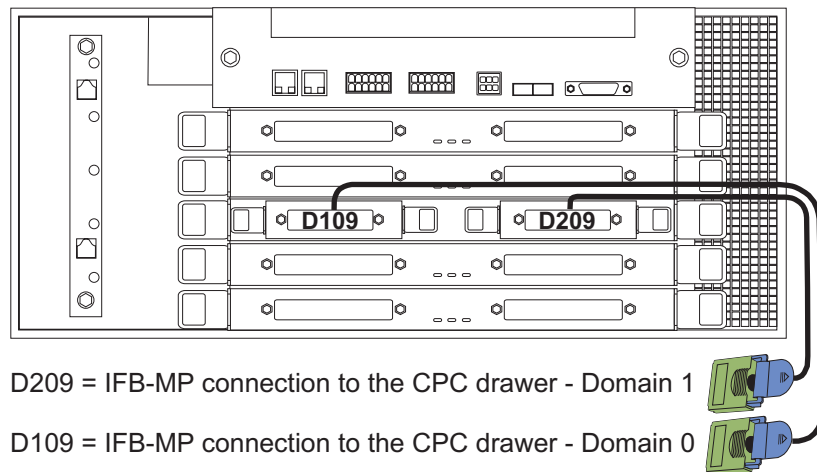
Card slot 07 - Domain 1

Card slot 08 - Domain 0

D109 D209

Card slot 10 - Domain 0

Card slot 11 - Domain 1



D209 = IFB-MP connection to the CPC drawer - Domain 1

D109 = IFB-MP connection to the CPC drawer - Domain 0

Figure 4. I/O drawer layout for the z10 BC

Legend:

IFB-MP

InfiniBand interface

D1 Half-high card in left side of slot

D2 Half-high card in right side of slot

For maximum availability of the device, OSA network, or coupling facility on z10 BC, you should consider the following guidelines:

- Choose channels plugged in different I/O domains (see Figure 4). An I/O domain contains four channel cards controlled by a single IFB-MP card. For example, the domain for the IFB-MP card in D109 contains slots 02, 05, 08, and 10.

Note: This is also recommended for optimum performance of your most heavily-used I/O devices.

When choosing the I/O domains to use, whether from different I/O drawers or the same I/O drawer, consider using a combination of I/O domains. When you must use IFB links from the CPC drawer, try to use IFB links from different HCA fanout cards. Refer to your PCHID report to determine which IFB links belong to which HCA fanout cards. If you have multiple paths to the device and multiple domains available that have the correct channel type, spreading the paths across as many HCAs as possible is also advisable.

Redundant I/O interconnect is a function that allows one IFB-MP back up another IFB-MP in case of a failure or repair. The IFB-MP card in slot 09 back up each other. Therefore, in the event of an InfiniBand cable or fanout card failure, the remaining IFB-MP card will control both domains. There are failures (for example, the IFB-MP card) that may prevent the redundant takeover, which is why it is advisable to spread your paths over multiple domains.

When configuring ICB-4 channel paths for the same target CPC or coupling facility, use different HCA fanout cards on the CPC drawer. This allows for continued connectivity if you have to run degraded after an MBA fanout card failure.

When configuring coupling over InfiniBand (CIB) channel paths for the same target CPC or coupling facility, use different HCA cards on the CPC drawer. This eliminates the HCA fanout card and the IFB cable as a single point of failure where all connectivity would be lost.

- If you define multiple paths from the same IFB link, distribute paths across different channel cards. Also, if you define multiple coupling facility channels to the same coupling facility or to the same ESA image, distribute paths across different coupling facility channel adapter cards or different coupling facility daughter cards.

z10 servers have eliminated the concept of FICON channels or channel cards having affinity to specific System Assist Processors (SAPs). With System z10 BC, each SAP handles FICON work on an on-demand basis. That is, as FICON work for any channel arrives, the next available SAP will handle that request. It does not matter if it is an outbound request or an inbound interrupt, the next available SAP will handle the FICON work. There is no requirement to configure for “SAP affinity,” because there is no affinity for FICON.

For the other channel types, the z10 servers automatically balance installed channel cards across all available SAPs. The processor attempts to assign an equal number of each channel card type to each available SAP. While all channels on a given I/O card are always in the same SAP, it is not predictable which I/O cards will be assigned to which SAPs. However, there are two exceptions. First, ICBs and Host Channel Adapters (HCAs) used for coupling are always given affinity to a SAP on the local book. Second, if an OSA channel is defined as OSD or a FICON channel is defined as FCP, these channels use QDIO architecture and, therefore, do not actually use any SAP resource during normal operations.

For all channel types, simply follow the preceding recommendations for configuring for RAS, and the SAPs will handle the workload appropriately.

InfiniBand and ICB-4 configurations

A z10 BC has one CPC drawer installed. The drawer can be configured with MBA fanout cards (for ICB-4 connections) and HCA fanout cards (for InfiniBand connections.)

The MBA fanout card is used to connect a z10 BC or z10 EC to a z9 EC, z9 BC, z990, or z890 server using an ICB-4 connection. (However, it can also connect two System z10 machines.) The MBA fanout card is designed to support two Self-Timed Interconnect (STIs) operating at a link rate of 2.0 GBps. The STI cable is 10 meters (33 feet) in length, of which 3 meters are reserved for intraserver connection.

There are four types of HCA fanout cards: HCA2-C, HCA2-O, HCA2-O LR, and HCA1-O. (HCA1-O is for z9 only.)

- The HCA2-C fanout card is used for internal connection to the I/O drawer only.
- The HCA2-O fanout card is used to connect a z10 BC to a z10 BC, z10 EC, or z9 at an distance of 150 meters (492 feet). The HCA2-O fanout attaches to a z10 BC. It is designed to support a two-port 12x IB-DDR optical link operating at a link rate of 6 GBps. The HCA1-O attaches to a z9. It is designed to support a two-port 12x IB-SDR optical link operating at a link rate of 3 GBps.
- The HCA2-O LR fanout card is used to connect a z10 server to a z10 server at an unrepeat distance of 10 km (6.2 miles). It is designed to support a two-port 1x IB-SDR or 1x IB-DDR optical link operating at a link rate of 2.5 Gbps (1x IB-SDR) or 5 gbps (1x IB-DDR).

The z10 BC is designed to contain up to a maximum of 12 total connections.

Figure 5 depicts the z10 BC IFB, ICB-4, and ISC-3 connections.

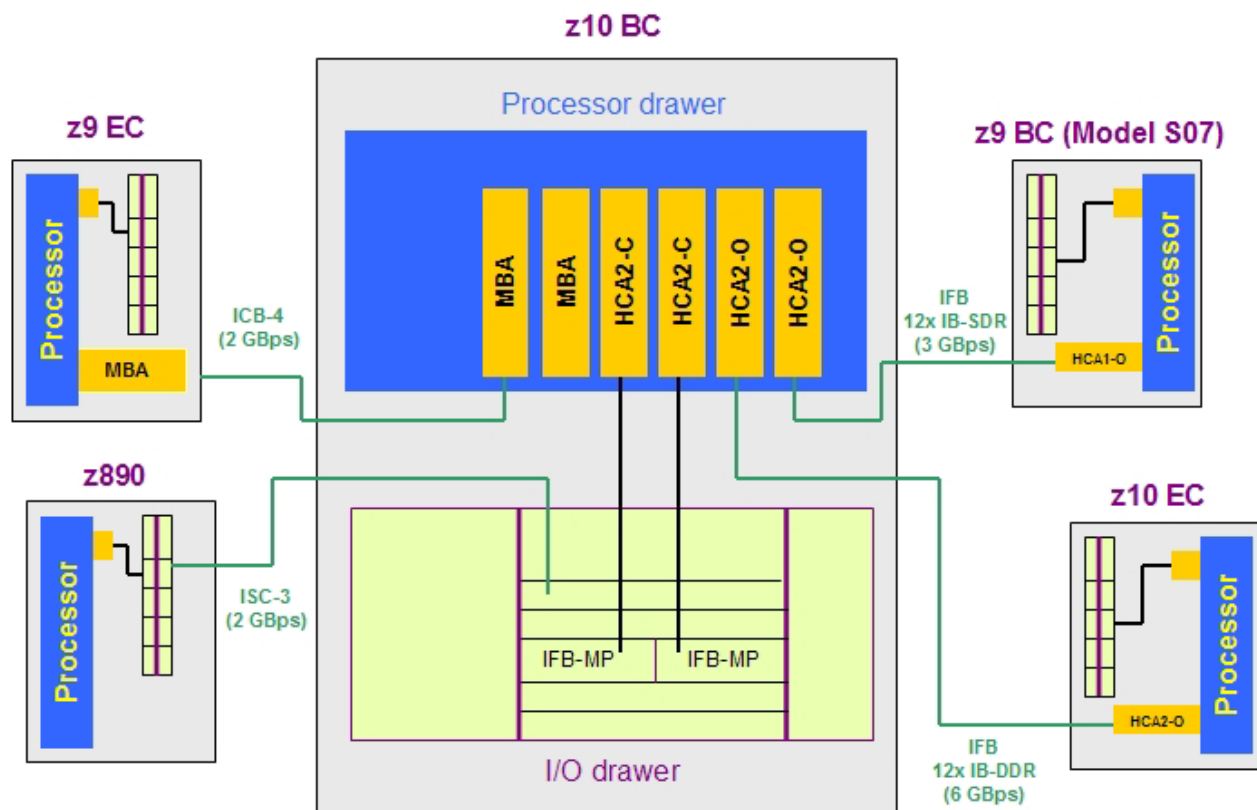


Figure 5. InfiniBand, ICB-4, and ISC-3 connections

Planning for channel subsystem

This section contains information to aid in the planning for maximum channel subsystem availability on z10 BC. It addresses ESCON, FICON, and OSA and ISC-3 channels; ICB-4, IFB, and IC links; and HyperSockets. The information shows the major components of the channel subsystem and suggests ways to configure the CSS for maximum availability.

The overall process of assigning CHPIDs, PCHIDs, and AIDs begins when you order the z10 BC or an MES to an existing machine. After placing the order, the configurator prepares a report (PCHID report) detailing the physical location for each channel in the machine. This report shows PCHID and AID assignments.

PCHID assignments

There are no default CHPIDs assigned. You are responsible for assigning the logical CHPID number to a physical location, identified by a PCHID number. You can complete this task in using either IOCP or HCD. The CHPID Mapping Tool may be used to help with these assignments. (Refer to “CHPID Mapping Tool” on page 40 for more information.)

You will use the data in the CFReport, which you can either obtain from your representative or retrieve from Resource Link, and the IOCP input for assigning PCHID values using the CHPID Mapping Tool.

A PCHID relates directly to a jack location (J01/J02) in a specific fanout card sot in a specific drawer. Table 8 shows the PCHID assignments for the jacks on MBA fanout cards (ICB-4 attachments) plugged into the drawer. HCAs are not assigned PCHID values. Each fanout slot is allocated 2 PCHID numbers. (Remember that slots D1/D2 and D9/DA do not have fanout cards plugged into them; therefore, they are not assigned PCHIDs.)

Table 8. PCHID assignments for MBA fanout cards

CPC Drawer	Slot number	Fanout card slots	PCHIDs
1	A21B	D3-D8	000-00F

Table 9 lists the PCHID assignments for slots in the I/O drawer. Only the active ports on an installed card are actually assigned a PCHID, the remainder are unused.

Except for ESCON sparing, the cards in the I/O drawer are assigned a PCHID starting with the first value in the range for the slot and drawer where the card is located. For ISC-3 cards, the first daughter is assigned the first two PCHID values of the slot. The second daughter is assigned the slot value plus 8 for the first port and plus 9 for the second port.

Table 9. PCHIDs assignments for I/O drawer

Slot	PCHID range			
	Drawer 4 A01B	Drawer 1 A06B	Drawer 2 A11B	Drawer 3 A16B
2	100 - 10F	180 - 18F	200 - 20F	280 - 28F
3	110 - 11F	190 - 19F	210 - 21F	290 - 29F
4	120 - 12F	1A0 - 1AF	220 - 22F	2A0 - 2AF
5	130 - 13F	1B0 - 1BF	230 - 23F	2B0 - 2BF
7	140 - 14F	1C0 - 1CF	240 - 24F	2C0 - 2CF
8	150 - 15F	1D0 - 1DF	250 - 25F	2D0 - 2DF
10	160 - 16F	1E0 - 1EF	260 - 26F	2E0 - 2EF
11	170 - 17F	1F0 - 1FF	270 - 27F	2F0 - 2FF

Direct drawer to drawer connections (z10 BC to z10 BC) or direct drawer to book connections (z10 BC to z10 EC, z9 EC, z9 BC, z990, or z890) can use ICB-4 links. (These links are used for connections to coupling facilities.) There can be up to 12 STI ports on a drawer, and a PCHID number is assigned to each one.

AID assignments

HCA2-O and HCA2-O LR fanout cards used for coupling are identified by adapter IDs (AIDs) rather than PCHIDs.

CHPID numbers need to be associated with ports on an adapter, and the AID is used for this purpose. You are responsible for assigning CHPIDs. You can use either IOCP or HCD. The CHPID assignment is done by associating the CHPID number to an AID and port. You cannot use the CHPID Mapping Tool to assign AID values.

You cannot change an AID assignment. After an AID is assigned, if an optical fanout card is moved on a z10 BC, the AID value moves with it.

There are two ports for each HCA. The CHPID to PCHID ratio was one-to-one, but an AID can have up to 16 associated CHPIDs.

Table 10 shows the initial AID assignments for the ports on the HCA fanout cards plugged into the CPC drawer.

Table 10. AID assignments for HCA fanout cards

CPC drawer number	Location	Fanout card slots	Possible AIDs
1	A21B	D3-D8	00-05

Each fanout slot is allocated one AID number. (Remember that slots D1/D2 and D9/DA do not have fanout cards plugged into them; therefore, they are not assigned AIDs.) For example, the allocation for the CPC drawer can be:

Fanout slot	AID
D3	00
D4	01
D5	02
D6	03
D7	04
D8	05

Note: These AID assignments can only be predicted for a new build machine. For an MES to an existing z10 Business Class, the PCHID report contains the AID assigned to each installed HCA and the AID that is assigned to any new HCA being installed. Therefore, if a new HCA is added to the drawer, the AID from the PCHID report should be used.

PCHID report

The PCHID report from the configurator provides details on the placement of all the I/O features in your order. Your representative will provide you with this report. Using this report and the guidelines listed in “Guidelines for maximum availability” on page 34, you can plan the configuration of your I/O.

Note: If you use the CHPID Mapping Tool to aid you in assigning PCHIDs to CHPIDs, the tool will provide you with a new report with your CHPID assignment in addition to the PCHID values.

Other resources available are the *System z Input/Output Configuration Program User's Guide for ICP IOCP* and the CHPID Mapping Tool. These resources are available on Resource Link.

CHPID Mapping Tool

The CHPID Mapping Tool is a Java-based standalone application available from IBM *Resource Link*, and it must be downloaded to your personal computer for use. Once downloaded, you can make CHPID assignments without further internet connections. As part of the CHPID Mapping Tool process, you will need a CFReport (which you can download from Resource Link or obtain from your representative) and an IOCP file.

Note: The CHPID Mapping Tool does not assign AID values.

The intent of the CHPID Mapping Tool is to ease installation of new z10 BC processors and for making changes to an already installed z10 BC processor either to make slight changes to the mapping or as part of an MES action to add or remove channel features on the processor.

The System z10 BC server **does not** have default CHPIDs assigned to ports as part of the initial configuration process. It is your responsibility to perform these assignments by using the HCD/IOCP definitions and optionally the CHPID Mapping Tool. The result of using the tool is an IOCP deck that will map the defined CHPIDs to the corresponding PCHIDs for your processor. However, there is no requirement to use the mapping tool. You can assign PCHIDs to CHPIDs directly in IOCP decks or through HCD, but it is much easier to use the tool to do the channel mapping and the tool can help make PCHID to CHPID assignments for availability.

For more information on the CHPID Mapping tool refer to any of the following:

- *System z CHPID Mapping Tool User's Guide*
- *CHPID Mapping Tool* on Resource Link.

Multiple Image Facility (MIF)

The Multiple Image Facility (MIF) allows channel sharing among multiple LPARs and optionally shares any associated I/O devices configured to these shared channels. MIF also provides a way to limit the logical partitions that can access a reconfigurable channel, spanned channel, or a shared channel to enhance security.

With multiple LCSSs, the CSS provides an independent set of I/O controls for each logical channel subsystem called a CSS image. Each logical partition is configured to a separate CSS image in order to allow the I/O activity associated with each logical partition to be processed independently as if each logical partition had a separate CSS. For example, each CSS image provides a separate channel image and associated channel path controls for each shared channel and separate subchannel images for each shared device that is configured to a shared channel.

With MIF, you can configure channels as follows:

- **ESCON** (TYPE=CNC, TYPE=CTC, TYPE=CVC, or TYPE=CBY), **FICON** (TYPE=FCV, TYPE=FC, or TYPE=FCP), **ISC-3 peer** (TYPE=CFP), **ICB-4 peer** (TYPE=CBP), **IC peer** (TYPE=ICP), **IFB peer** (TYPE=CIB), **HiperSockets** (TYPE=IQD), and **OSA** (TYPE=OSC, TYPE=OSD, TYPE=OSE or TYPE=OSN).

You can configure a channel path as:

- An unshared dedicated channel path to a single LPAR.

- An unshared reconfigurable channel path that can be configured to only one logical partition at a time it can be moved to another logical partition within the same LCSS.
- A shared channel path that can be concurrently used by the ESA/390 images or CF logical partitions within the same LCSS to which it is configured.

With MIF and multiple channel subsystems, shared and spanned channel paths can provide extensive control unit and I/O device sharing. MIF allows all, some, or none of the control units attached to channels to be shared by multiple logical partitions and multiple CSSs. Sharing can be limited by the access and candidate list controls at the CHPID level and then can be further limited by controls at the I/O device level.

For example, if a control unit allows attachment to multiple channels (as is possible with a 3990 control unit), then it can be shared by multiple logical partitions using one or more common shared channels or unique unshared channel paths.

Spanned channels

With multiple LCSSs, transparent sharing of internal (ICs and HiperSockets) and external (FICON, ICB-4, ISC-3, IFB, OSA) channels across LCSSs is introduced, extending Multiple Image Facility (MIF). MIF allows sharing of channel resources across LPARs. ICs, HiperSockets, FICON (except when defined as FCV channels), ICB-4s, ISC-3s, IFBs, and OSA features can all be configured as MIF spanning channels.

Spanning channels is the ability for the channels to be configured to multiple LCSSs, and be transparently shared by any or all of the configured LPARs without regard to the Logical Channel Subsystem to which the partition is configured. For information on the channel CHPID types and spanning capabilities, refer to Table 7 on page 32.

You can configure the following as a spanned channel:

- **FICON (TYPE=FC or TYPE=FCP), ISC-3 peer (TYPE=CFP), ICB-4 peer (TYPE=CBP), IC peer (TYPE=ICP), IFB peer (TYPE=CIB), HiperSockets (TYPE=IQD), and OSA (TYPE=OSC, TYPE=OSD, TYPE=OSE or TYPE=OSN)**

They can be shared by LPARs in different logical channel subsystems.

Internal coupling and HiperSockets channels

Internal coupling (IC) channels and HiperSockets are virtual attachments and, as such, require no real hardware. However, they do require CHPID numbers and they do need to be defined in the IOCDs. The CHPID type for IC channels is ICP; the CHPID type for HiperSockets is IQD.

- It is suggested that you define a minimum number of ICP CHPIDs for Internal Coupling. For most customers, IBM suggests defining just one ICP for each coupling facility (CF) LPAR in your configuration. For instance, if your z10 BC configuration has several ESA LPARs and one CF LP, you would define one pair of connected ICP CHPIDs shared by all the LPARs in your configuration. If your configuration has several ESA LPARs and two CF logical partitions, you still would define one connected pair of ICP CHPIDs, but one ICP should be defined as shared by the ESA images and one of the CF LPARs, while the other ICP is defined as shared by the ESA LPARs and the other CF LPAR. Both of these examples best exploit the peer capabilities of these coupling channels by using the “sending” and “receiving” buffers of both channels. If your ESA images and

CF images are in different CSSs and you want to exploit the optimal use of ICP then your ICP CHPIDs must be defined as spanned.

- Each IQD CHPID represents one internal LAN. If you have no requirement to separate LAN traffic between your applications, only one IQD CHPID needs to be defined in the configuration. If the partitions sharing the LAN are in different LCSSs your IQD CHPID must be defined as spanned.

IOCP considerations

ICP IOCP supports the z10 BC model CPCs and multiple LCSSs. Refer to *System z Input/Output Configuration Program User's Guide for ICP IOCP* for more information.

IOCP allows you to define controls for multiple channel subsystems. This includes changes to the way you define LPARs, channel paths, and I/O devices.

LPAR definition

Use the RESOURCE statement to define LCSSs and the logical partitions in each LCSS. You can also assign a MIF image ID to each LPAR. If you do not specify a MIF image ID using the RESOURCE statement, ICP IOCP assigns them. Any LPARs not defined will be reserved and available to be configured later using dynamic I/O.

Channel path definition

You can define shared channel paths in addition to dedicated and reconfigurable channel paths. The CHPID statement has an additional SHARED keyword to accomplish this. You can also define spanned channel paths using the PATH keyword. You can define:

- All channel paths as dedicated or reconfigurable.
- Only CNC, CTC, FCV, FC, FCP, CFP, CBP, ICP, IQD, CIB, OSC, OSD, OSE and OSN channel paths as shared.
- Only FC, FCP, CFP, CBP, ICP, IQD, CIB, OSC, OSD, OSE, and OSN channel paths as spanned.

ICP IOCP provides access controls for spanned, shared or reconfigurable channel paths. Parameters on the PART | PARTITION or NOTPART keyword on the CHPID statement allow you to specify an access list and a candidate list for spanned, shared and reconfigurable channel paths.

The access list parameter specifies the logical partition or logical partitions that will have the channel path configured online at logical partition activation following the initial power-on reset of an LPAR IOCDS. For exceptions, refer to *System z10 Processor Resource/Systems Manager Planning Guide*.

The candidate list parameter specifies the LPARs that can configure the channel path online. It also provides security control by limiting the logical partitions that can access shared or reconfigurable channel paths.

Note: PR/SM LPAR manages the channel path configuration across POR. Refer to *System z10 Processor Resource/Systems Manager Planning Guide*.

I/O device definition

You can specify either the optional PART | PARTITION keyword or the optional NOTPART keyword on the IODEVICE statement to limit device access by logical partitions for devices assigned to shared ESCON, FICON, or OSA channels, or HiperSockets. (The IODEVICE candidate list is not supported for shared CFP, CBP, CIB, or ICP CHPIDs.)

By limiting access to a subset of logical partitions, you can:

- Provide partitioning at the device level.
- Provide security at the device level.
- Better manage the establishment of logical paths.

Hardware Configuration Definition (HCD) considerations

HCD provides the capability to make both dynamic hardware and software I/O configuration changes. It also provides:

- An online, interactive way to more useably manage the I/O configuration than IOCP.
- The capability to define the I/O configuration for dynamic or nondynamic I/O configuration purposes.

HCD allows you to define LPAR controls for defining LPARs, channel paths, and I/O devices. The following HCD panels (or corresponding HCM dialogs) support these controls.

Add Partition

Allows explicit definition of LPARs and associated LPAR numbers.

Define Access List

Allows definition of initial access list for channel path access control of shared and reconfigurable channel paths.

Define Candidate List (for channel paths)

Allows definition of candidate list for channel path access control of shared and reconfigurable channel paths.

Define Candidate List (for devices)

Allows definition of candidate list for device access control for devices assigned to shared channels.

Add Processor

Allows you to determine the capabilities of a CPC.

Add Channel Path

Operation mode field allows definition of a channel path as dedicated, reconfigurable, or shared.

Define Device / Processor

Additional field to specify candidate list.

Chapter 5. I/O connectivity

This chapter discusses the channels associated with the z10 BC I/O connectivity. You can also refer to Table 4 on page 19 for a summary of the I/O channel characteristics.

FICON and FCP channels

The FICON Express channel uses the industry standard Fibre Channel Standard as a base. It is an upper layer protocol that maps the channel architecture on the general transport vehicle used throughout the industry for such other upper layer protocols as SCSI, IPI, and IP, among others. This transport vehicle includes the physical definition, the transmission protocol, and signalling protocol that is the same for all of the other upper layer protocols.

The FICON Express8, FICON Express4, FICON Express2, and FICON Express features are designed to provide connectivity to servers, switches, directors, disk, tape, and printers. There are two CHPID types that can be specified using IOCP or HCD. Each channel has its own unique CHPID type:

- CHPID type FC — native FICON, High Performance FICON for System z (zHPF), and channel-to-channel (CTC)
- CHPID type FCP — Fibre Channel Protocol (FCP) for communication with SCSI devices

The FICON Express LX feature supports a third CHPID type:

- CHPID type FCV — Fibre Channel Converted to communicate with ESCON control units using the FICON Bridge card in the ESCON Director Model 5 (9032-005), which was withdrawn from marketing December 31, 2004. FICON Bridge was designed as a migration aid; to facilitate a transition from ESCON to FICON.

FICON builds upon the strengths of ESCON. The FICON implementation enables full duplex data transfer. So data travels both directions simultaneously, rather than the ESCON half duplex data transfer. Furthermore, concurrent I/Os can occur on a single FICON channel, a fundamental difference between FICON and ESCON. The data rate droop is minimal with FICON even at distances up to 100 km.

Native FICON supports up to 32 concurrent I/O operations per second. ESCON supports one I/O operation at a time. FICON bridge supports up to eight I/O operations at a time to different control units.

In conjunction with the Fibre Channel protocol (FCP), N_Port ID Virtualization (NPIV) is supported, which allows the sharing of a single physical FCP channel among operating system images.

FICON Express8 features

The FICON Express8 features conform to the Fibre Connection (FICON) architecture, the High Performance FICON on System z (zHPF) architecture, and the Fibre Channel Protocol (FCP) architecture, providing connectivity between any combination of servers, directors, switches, and devices (control units, disks, tapes, printers) in a Storage Area Network (SAN). FICON Express8 provides increased performance over FICON Express4, FICON Express2, or FICON Express.

There are two FICON Express8 features for the IBM System z10 BC — FICON Express8 10KM LX and FICON Express8 SX.

Each feature has four channels per feature. Each of the four independent channels supports a link rate of 2 gigabits (Gbps), 4 Gbps, or 8 Gbps per second, depending upon the capability of the attached switch or device, with autonegotiation to 2 Gbps or 4 Gbps for existing devices. A link data rate of 1 Gbps is not supported. Each channel utilizes small form factor pluggable optics (SFPs) with LC duplex connectors. The optics allow each channel to be individually repaired without affecting the other channels.

Each FICON Express8 feature supports 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 applications, GDPS®, and remote copy.

FICON Express8 10KM LX (FC 3325)

All the channels on a single FICON Express8 10KM LX feature are the same type, 10KM LX. FICON Express8 10KM LX utilizes a long wavelength (LX) laser as the optical transceiver and supports use of a 9 micron single mode fiber optic cable terminated with an LC duplex connector.

FICON Express8 10KM LX supports unrepeat distances up to 10 km (6.2 miles).

FICON Express8 10KM LX (CHPID type FC or FCP) can be defined as a spanned channel and can be shared among LPARs within and across LCSS.

FICON Express8 SX (FC 3326)

All the channels on a single FICON Express8 SX feature are the same type, SX. FICON Express8 SX utilizes a short wavelength (SX) laser as the optical transceiver and supports use of a 50 micron multimode fiber optic cable or a 62.5 micron multimode fiber optic cable terminated with an LC duplex connector.

For details about the unrepeat distances for FICON Express8 SX, refer to *System z Planning for Fiber Optic Links*.

FICON Express4 features

The FICON Express4 feature conforms to the Fibre Connection (FICON) architecture, the High Performance FICON for System z (zHPF) architecture, and the Fibre Channel (FC) architecture, providing connectivity between any combination of servers, directors, switches, disks, tapes, printers, and devices in a Storage Area Network (SAN).

The FICON Express4 features for the IBM System z10 BC include:

- **FICON Express4 10KM LX (FC 3321)**

FICON Express4 10KM LX has four channels per feature, is designed to support unrepeat distances up to 10 km (6.2 miles) over single mode fiber optic cabling.

- **FICON Express4 4KM LX (FC 3324)**

FICON Express4 4KM LX has four channels per feature, is designed to support unrepeat distances up to 4 km (2.5 miles) over single mode fiber optic cabling.

- **FICON Express4 SX (FC 3322)**

FICON Express4 SX has four channels per feature, is designed to carry traffic over multimode fiber optic cabling.

- **FICON Express4-2C 4KM LX (FC 3323)**

FICON Express4-2C 4KM LX has two channels per feature, is designed to support unrepeatable distances up to 4 km (2.5 miles) over single mode fiber optic cabling.

- **FICON Express4-2C SX (FC 3318)**

FICON Express4-2C SX has two channels per feature, is designed to carry traffic over multimode fiber optic cabling.

All channels on a single FICON Express4 feature are of the same type: 4KM LX, 10KM LX, or SX. You can carry forward your current FICON Express2 LX, FICON Express2 SX, FICON Express LX, and FICON Express SX features from z9 BC or z890 to System z10 BC.

FICON Express4 supports a 4 Gbps link data rate with auto-negotiation to 1, 2, or 4 Gbps for synergy with existing switches, directors, and storage devices. An entry level 4KM LX feature supporting two channels per feature for data centers with limited requirements for single mode fiber optic cabling connectivity is offered.

Note: You need to ensure that the tactical as well as the strategic requirements for your data center, Storage Area Network (SAN), and Network Attached Storage (NAS) infrastructures are taken into consideration as you employ 2 Gbps and beyond link data rates.

Mode Conditioning Patch (MCP) cables are only supported at the 1 Gbps link data rate.

Effective October 27, 2009, FICON Express4 is withdrawn from marketing. After that time, FICON Express4 cannot be ordered. FICON Express8 replaces FICON Express4.

FICON Express2 features

The FICON Express2 LX (FC 3319) and FICON Express2 SX (FC 3320) features can only be carried forward to System z10 BC. They are not orderable.

The FICON Express2 feature conforms to the Fibre Connection (FICON) architecture, the High Performance FICON for System z (zHPF) architecture, and the Fibre Channel (FC) architecture, providing connectivity between any combination of servers, directors, switches, and devices in a Storage Area Network (SAN). FICON Express2 supports cascading to concentrate and minimize the number of cross-site connections and reduce implementation costs for disaster recovery applications, Graphically Dispersed Parallel Sysplex (GDPS), and remote copy.

FICON Express features

The FICON Express LX (FC 2319) and FICON Express SX (FC 2320) features can only be carried forward to System z10 BC. They are not orderable.

FICON Express LX can be defined as CHPID type FCV to allow communication with ESCON control units using the ESCON Director Model 5 with the bridge feature. However, migration to native FICON is recommended.

Name server registration

Registration information is provided on the name server for both FICON and FCP, which enhances problem determination, analysis, and manageability of the storage area network (SAN).

High Performance FICON for System z (zHPF)

Enhancements (called High Performance FICON for System z (zHPF)) have been made to the FICON architecture for protocol simplification and efficiency. With these enhancements, FICON Express8, FICON Express4 and FICON Express2 features will see a performance improvement in small block (4K) data transfers.

zHPF is supported by z/OS. zHPF applies to all FICON Express8, FICON Express4 and FICON Express2 features (CHPID type FC) and is exclusive to System z10.

Improved performance at extended distance

Enhancements have been made to the industry standard FICON architecture (FC-SB-3) to avoid performance degradation at extended distances. This enhancement is accomplished by implementing a new protocol for Information Unit (IU) pacing. This enhancement allows the channel to “remember” the last pacing update for use on subsequent operation to help avoid degradation of performance at the start of each new operation.

To exploit this enhancement, the control unit must also support the new pacing protocol.

The MIDAW facility

The Modified Indirect Data Address Word (MIDAW) facility is a system architecture and software designed to improve FICON performance. This facility is exploited by z/OS access methods and supported by z/VM 5.3 or later for guests exploitation. The MIDAW facility can:

- Significantly improve FICON performance for extended format data sets. Non-extended data sets can also benefit from MIDAW.
- Improve channel utilization and can significantly improve I/O response time by reducing FICON channel and control unit overhead.

Multipath Initial Program Load (IPL)

If I/O errors occur during the IPL, z/OS on System z10 BC allows the system to attempt an IPL on alternate paths, if the paths are available. The system will attempt the IPL on an alternate path until all paths have been attempted or until the IPL is successful.

This function is applicable for all FICON features with CHPID type FC.

Purge path extended

The purge path extended function provides enhanced capability for FICON problem determination. The FICON purge path error-recovery function is extended to transfer error-related data and statistics between the channel and entry switch, and from the control unit and its entry switch to the host operating system.

FICON purge path extended is supported by z/OS and z/OS.e. FICON purge path extended applies to the FICON features when configured as a native FICON channel.

Fibre channel analysis

You can use the Fibre Channel Analyzer task on the HMC to identify fiber optic cabling issues in your Storage Area Network (SAN) fabric without contacting IBM service personnel. All FICON channel error information is forwarded to the HMC where it is analyzed to help detect and report the trends and thresholds for all FICON channels on System z10. This report shows an aggregate view of the data and can span multiple systems.

This applies to FICON channels exclusively (CHPID type FC).

Fibre Channel Protocol (FCP) for SCSI devices

Fibre Channel (FC) is a computer communications protocol that attempts to combine the benefits of both channel and network technologies. Fibre Channel made the biggest impact in the storage arena, specifically using Small Computer System Interface (SCSI) as an upper layer protocol.

Fibre Channel is broken down into five layers: FC-0, FC-1, FC-2, FC-3, and FC-4. The layers define the following functions:

- **FC-0** defines the physical characteristics
- **FC-1** defines the character encoding and link maintenance
- **FC-2** defines the frame format, flow control, classes of service
- **FC-3** defines the common services

FICON and FCP implement those layers, unchanged.

- **FC-4** defines the upper layer protocol mapping which includes SCSI as well as Fibre Channel - Single Byte-2 (FC-SB-2), which is FICON.

The Fibre Channel Protocol (FCP) capability, supporting attachment to Small Computer Systems Interface (SCSI) is based on the Fibre Channel (FC) standards defined by the INCITS, and published as ANSI standards. SCSI devices in Linux on System z environments, as well as SCSI devices defined to z/VM and z/VSE, are based on the Fibre Channel standards. FC is an upper layer fibre channel mapping of SCSI on a common stack of Fibre Channel physical and logical communication layers. HIPPI, IPI, IP, and FICON (FC-SB-2) are other examples of upper layer protocols.

SCSI is an industry-standard protocol that is supported by a wide range of controllers and devices that complement the classical System z9 and System z10 storage attachment capability through FICON and ESCON channels. FCP channels on System z9 and System z10 are provided to enable operating systems on System z9 and System z10 to access industry-standard SCSI storage controllers and devices.

The System z9 and System z10 FCP function can facilitate the consolidation of UNIX server farms onto System z9 and System z10, protecting investments in SCSI-based storage.

FCP is the base for open industry-standard Fibre Channel networks or Storage Area Networks (SANs).

Fibre Channel networks consist of servers and storage controllers and devices as end nodes, interconnected by Fibre Channel switches, directors, and hubs. While switches and directors are used to build Fibre Channel networks or fabrics, Fibre Channel loops can be constructed using Fibre Channel hubs. In addition, different

types of bridges and routers may be used to connect devices with different interfaces (like parallel SCSI). All of these interconnects may be combined in the same network.

For information about the configurations supported by the System z9 and System z10 FCP channel, refer to “Configurations.”

An FCP channel is defined in the IOCP as channel type FCP and is available on FICON features.

FCP channels support full-fabric support. The FCP full-fabric support means that multiple numbers of directors/switches can be placed between the server and FCP/SCSI device, thereby allowing many hops through a storage network for I/O connectivity.

In addition, for FCP channels, a high integrity fabric solution is not required but is recommended. If an FCP Interswitch Link (ISL) is moved, data could potentially be sent to the wrong destination without notification.

The FICON Express, FICON Express2, FICON Express4, and FICON Express8 features, when defined as CHPID type FCP in the IOCP, support storage controllers and devices with an FCP interface in z/VM, z/VSE, and Linux on System z environments.

Each port on a single FICON card can be configured individually and can be a different CHPID type.

Configurations

Storage controllers and devices with an FCP interface can be directly attached to the System z9 or System z10 server (point-to-point connection), or by using Fibre Channel switches or directors. A storage controller or device with an appropriate FCP interface may be attached to each port of a FICON feature, or of a Fibre Channel switch or director.

In addition, the following devices and controllers can be attached to each port on a Fibre Channel switch or director:

- **FC-AL controllers or devices, and FC-AL hubs**

If the switch or director supports the Fibre Channel Arbitrated Loop (FC-AL) protocol, devices implementing this protocol may be attached to that port and accessed from the System z9 or System z10 server. Devices typically implementing the FC-AL protocol are tape units and libraries, and low-end disk controllers.

If the switch or director does not support the FC-AL protocol, you can also install a FC-AL bridge between the switch or director and the FC-AL controller or device.

If more than one FC-AL controller or device should be attached to a FC-AL switch port, it is convenient to use a Fibre Channel hub, where multiple devices with a FC-AL interface can be directly attached to that hub.

- **Fibre-Channel-to-SCSI bridges**

Fibre-Channel-to-SCSI bridges can be used to attach storage controllers and devices implementing the electrical, parallel SCSI interface. Different types of Fibre-Channel-to-SCSI bridges may support different variants of the parallel SCSI interface, such as Low Voltage Differential (LVD), High Voltage Differential (HVD), Single Ended, wide (16-bit) versus narrow (8-bit) interfaces, and different link speeds.

Each System z9 or System z10 FCP channel (CHPID) can support up to 480 subchannels, where each subchannel represents a communication path between software and the FCP channel. Refer to “Channel and device sharing” on page 52 for more information.

Host operating systems sharing access to a System z9 or System z10 FCP channel can establish in total up to 2048 concurrent connections to up to 510 different remote Fibre Channel ports associated with Fibre Channel controllers.

The total number of concurrent connections to end devices, identified by logical unit numbers (LUNs), must not exceed 4096.

I/O devices

The FCP channel implements the FCP standard as defined by the INCITS Fibre Channel Protocol for SCSI (FCP), and Fibre Channel Protocol for SCSI, Second Version (FCP-2), as well as the relevant protocols for the SCSI-2 and SCSI-3 protocol suites. Theoretically, each device conforming to these interfaces should work when attached to a System z9 or System z10 FCP channel as previously defined. However, experience tells us that there are small deviations in the implementations of these protocols. Therefore, it is advisable to do appropriate conformance and interoperability testing to verify that a particular storage controller or device can be attached to an FCP channel in a particular configuration (i.e. attached via a particular type of Fibre Channel switch, director, hub, or Fibre-Channel-to-SCSI bridge).

Also, for certain types of FCP and SCSI controllers and devices, specific drivers in the operating system may be required in order to exploit all the capabilities of the controller or device, or to cope with unique characteristics or deficiencies of the device.

For a list of storage controllers and devices that have been verified to work in a Fibre Channel network attached to a System z9 or System z10 FCP channel, and for specific software requirements to support FCP and SCSI controllers or devices, refer to <http://www.ibm.com/servers/eserver/zseries/connectivity>.

Addressing

FCP channels use the Queued Direct Input/Output (QDIO) architecture for communication with the operating system. IOCP is only used to define the QDIO data devices. The QDIO architecture for FCP channels, derived from the QDIO architecture that had been defined for communications via an OSA card, defines data devices that represent QDIO queue pairs, consisting of a request queue and a response queue. Each queue pair represents a communication path between an operating system and the FCP channel. It allows an operating system to send FCP requests to the FCP channel via the request queue. The FCP channel uses the response queue to pass completion indications and unsolicited status indications to the operating system.

IOCP is not used to define the actual Fibre Channel storage controllers and devices, nor the Fibre Channel interconnect units such as switches, directors, or bridges. IOCP is only used to define the QDIO data devices. The Fibre Channel devices (end nodes) in a Fibre Channel network are addressed using World Wide Names (WWNs), Fibre Channel Identifiers (IDs), and Logical Unit Numbers (LUNs). These addresses are configured on an operating system level, and passed to the FCP channel together with the corresponding Fibre Channel I/O or service request via a logical QDIO device (queue).

Channel and device sharing

An FCP channel can be shared between multiple operating systems, running in a logical partition or as a guest operating system under z/VM. Under z/VM, multiple z/VM, CMS, Linux on System z, and z/VSE guests are able to share SCSI channels and devices using z/VM Fixed Block Architecture (FBA) emulation. To access the FCP channel directly, each operating system needs its own QDIO queue pair, defined as a data device on an FCP channel directly in the IOCP.

Each FCP channel can support up to 480 QDIO queue pairs. This allows each FCP channel to be shared among 480 operating system instances.

Channel and device sharing using NPIV: N_Port ID Virtualization (NPIV) allows the sharing of a single physical FCP channel and attached devices, logical units, among operating system images, whether in logical partitions or as z/VM guests in virtual machines. This is achieved by assigning a unique WWPN to each subchannel that is defined for an FCP Channel using IOCP.

Each operating system instance using such a subchannel and its associated QDIO queues therefore also uses its own WWPN. When the operating system image starts using its subchannel, the FCP channel performs a login to the Fibre Channel fabric and acquires a unique Fibre Channel ID, also called N_Port ID. This ID is used in all further Fibre Channel communication that is done on behalf of this operating system image. System z9 and System z10 firmware takes care for the persistent and unique assignment of WWPNS to FCP subchannel.

Access controls based on the assigned WWPN can be applied in the SAN environment, using standard mechanisms such as zoning in FC switches and Logical Unit Number (LUN) masking in the storage controllers. You can configure the SAN prior to the installation of a new machine using the WWPN prediction tool available on Resource Link.

NPIV exploitation requires a Fibre Channel director or switch that supports the NPIV standard. If such a director or switch is installed, NPIV mode can be enabled for the FCP channel that attaches to this Fibre Channel switch or director through the Support Element. This enablement can be done on logical partition base, i.e., per FCP channel image.

NPIV is not supported in a point-to-point topology.

Channel and device sharing without NPIV: Without NPIV support, multiple operating system images can still concurrently access the same remote Fibre Channel port through a single FCP channel. However, Fibre Channel devices or logical units, identified by their LUNs, cannot be shared among multiple operating system images through the same FCP channel.

Positioning

FCP and SCSI are industry-standard protocols, which have been implemented by many vendors in a large number of different types of storage controllers and devices. These controllers and devices have been widely accepted in the market place and proven to be adequate to meet the requirements regarding reliability, availability, and serviceability (RAS) in many environments.

However, it must be noted that there are some advanced, unique RAS characteristics of System z9 and System z10 storage attachments based on ESCON and FICON attachments, using System z9 and System z10 channel programs (and the Extended Count Key Data (ECKD™) protocol in the case of disk

control units), that may not be readily available in such an FCP or SCSI based world. Therefore, whenever there are very stringent requirements regarding isolation, reliability, availability, and serviceability, a conscious decision must be made whether FCP attached storage controllers and devices or classical System z9 and System z10 FICON or ESCON attached control units should be used. Customers requiring the more robust RAS characteristics should choose FICON or ESCON channels.

SCSI Initial Program Load (IPL)

This function allows you to IPL an operating system from an FCP-attached disk, to execute either in a logical partition or as a guest operating system under z/VM. In particular, SCSI IPL can directly IPL a System z10 BC operating system that has previously been installed on a SCSI disk. Thus, there is no need for a classical channel (ESCON or FICON) attached device, such as an ECKD disk control unit, in order to install and IPL a System z10 BC operating system. The IPL device is identified by its Storage Area Network (SAN) address, consisting of the WWPN of the disk controller and the Logical Unit Number (LUN) of the IPL device.

You can also IPL a standalone-dump program from an FCP channel attached SCSI disk. The standalone-dump program can also store the generated dump data on such a disk.

SCSI IPL in z/VM allows Linux on System z, z/VSE, and other guest operating systems that support SCSI IPL to be IPLed from FCP-attached SCSI disk, when z/VM is running on a System z10 BC. Therefore, z/VM, z/VSE, and Linux on System z guests may be started and run completely from FCP channel attached disk in your hardware configuration.

z/VM provides the capability to install z/VM from a DVD to an Enterprise Storage Server® (ESS) SCSI disk emulated as a Fixed Block Architecture (FBA) disk as well as an Enterprise Storage Server from a DVD to a 3390 disk. Thus, z/VM and its Linux on System z guests may be started and run completely from FCP disks on your hardware configuration. Refer to z/VM subset of the 2098DEVICE Preventive Service Planning (PSP) bucket for any service required for z/VM support for SCSI IPL.

z/VM supports SCSI-attached disks to be used for installation, IPL, and operations such as storing dumps, and other functions, while continuing to provide support for ESCON-attached or FICON-attached disk or tape.

z/VM SCSI support allows a Linux on System z server farm and z/VSE to be deployed on z/VM in a configuration that includes only SCSI disks.

z/VM provides the capability to dump Linux on System z guests to FCP-attached SCSI disks. Benefits include:

- More guest virtual memory can be dumped because SCSI disks can be larger than ECKD disks
- Avoids the need to convert a VMDUMP into a Linux tool format
- Allows the same SCSI dump mechanisms to be used when running Linux for System z in an LPAR and in a z/VM virtual machine.

For Linux on System z support for SCSI IPL, refer to this website:
<http://www.ibm.com/developerworks/linux/linux390/>.

For additional information on:

- How to use SCSI IPL for a logical partition, refer to the *System z10 Support Element Operations Guide* or to the *System z Hardware Management Console Operations Guide*
- Messages that can show up on the operating systems console on the SE or Hardware Management Console, refer to *System z Small Computer Systems (SCSI) IPL - Machine Loader Messages*
- How to use SCSI IPL for a z/VM guest, refer to <http://www.vm.ibm.com/pubs> for appropriate z/VM publications
- How to prepare a Linux on System z IPL disk or a Linux on System z dump disk, refer to <http://www.ibm.com/developerworks/linux/linux390/> for appropriate Linux on System z publications.

ESCON channels

The ESCON channel provides a 17 MBps link data rate between host and control units for I/O devices. ESCON supports half-duplex data transfers over 62.5 multimode fiber optic cabling.

The ESCON channel provides a light-emitting diode (LED) light source for fiber optic cables. It can extend up to 3 kilometers (1.86 U.S. miles), a range that can be further extended to 6 or 9 kilometers (km) by retransmission through one or two ESCON directors.

With the availability of two LCSSs, you can define a maximum of 512 ESCON channels on your z10 BC, up to a maximum of 28 features per system. The maximum number of configurable channels is 256 per LCSS and per operating system image. The high density ESCON feature has 16 ports, 15 of which can be activated for your use. One port is always reserved as a spare, in the event of a failure of one of the other ports. When four ports are ordered, two 16-port ESCON features are installed and two ports are activated on each feature. After the first pair, ESCON features are installed in increments of one. ESCON channels continue to be ordered in increments of four channels.

ESCON supports these operating system environments: z/OS, z/VM, z/VSE, TPF, and Linux on System z.

ESCON channels affect the performance of the channel subsystem. Maximizing channel subsystem performance is an important consideration in configuring I/O devices to a z10 BC general purpose model CPC. Channel subsystem performance depends on the factors described in this chapter.

For an explanation of basic ESCON channel concepts, refer to *Introducing Enterprise Systems Connection*. For detailed information about synchronous and nonsynchronous I/O operation, refer to *Storage Subsystem Library Introduction to Nonsynchronous Direct Access Storage Subsystems*.

ESCON converter operation

You can configure ESCON converter channels (attached to a parallel converter - for example, the IBM 9034 or the Optica 34600 FXBT) for block and byte multiplexer mode of operation.

As many as eight channel paths are available to attach to any I/O device. During any I/O operation, one of the available channel paths is selected. Channel path selection is a hardware function rather than a function of the system control program.

At the start of an I/O operation, a central processor signals the channel subsystem that an I/O operation is needed. An I/O request is posted to a queue; meanwhile, instruction execution in the central processor continues.

Channel multiplexing modes

The data mode of operation is determined by the multiplexer mode (block or byte). This is selected for specific channels when either the CPC or a logical partition is initialized.

Block Multiplexer Mode of Operation: In block multiplexer mode of operation, a device stays connected to a channel continuously during the transfer of a full block of data.

Block multiplexer mode of operation allows a control unit to present “channel end” and to disconnect from a channel at the completion of a specified operation. “Device End” is presented at a later point. During the interval between “channel end” and “device end” another device attached to the same channel can be started or can complete an operation that is ready. However, if the second device does connect to the same channel during this interval, the first device may find the channel busy when it tries to reconnect, and then the first device must wait for service.

ESCON can be configured for block multiplexer mode of operation. In block multiplexer mode of operation, ESCON channels configured as CVC channel paths can operate in either interlock (high-speed transfer) mode or in data-streaming mode. They can also be attached to control units that operate in high-speed transfer or in data-streaming mode. Data rates can be as high 4.5 MBps for ESCON CVC channel paths.

Byte multiplexer mode of operation: Byte interleave mode of operation allows the execution of multiple I/O operations concurrently. Byte multiplexer mode permits several relatively slow-speed I/O devices to operate at the same time. Each addressed device requesting service is selected for transfer of a byte or a group of bytes to or from main storage. Bytes from multiple devices are interleaved on the channel and routed to or from the desired locations in main storage.

The load that a byte multiplexer channel can sustain is variable. It is governed by I/O device performance factors such as the data transfer rate, device buffers, number of bytes per data burst on the channel, channel program requirements, synchronized mechanical motion, and priority sequence position on the I/O interface.

ESCON converter channels (defined as CBY) can be configured for byte multiplexer mode of operation. In byte multiplexer mode of operation, ESCON channels configured as CBY channel paths can operate in either byte multiplexer mode or in burst mode. CBY channels require a 9034 ESCON converter. Byte multiplexer mode permits several relatively slow-speed I/O devices to operate at the same time.

Refer to the 2098IO subset id in the 2098DEVICE upgrade ID of the preventive service planning (PSP) bucket for prerequisite 9034 EC level information.

Byte multiplexer mode and burst mode: A byte multiplexer channel can be monopolized by one I/O device (burst mode) or shared by many I/O devices (byte multiplexer mode). The number of bytes transferred at a time in byte multiplexer mode can be one (single byte transfers) or more than one (multibyte transfers). Most control units that operate in byte multiplexer mode can also operate in burst

mode. A manually set switch at the control unit determines whether the control unit operates in burst mode or byte multiplexer mode.

Some devices offer a choice of how many bytes are transferred during a single data transfer sequence in byte multiplexer mode.

Because most of the time spent in a data-transfer control sequence is for control, increasing the burst size (the number of bytes transferred per sequence) results in a relatively small increase in the total channel busy time for the sequence. Also, increasing the burst size reduces the number of data transfer sequences required. The net effect is a significant improvement in channel efficiency and a higher allowable data rate.

Burst mode, although most effective in the use of channel resources, can cause another device on the byte multiplexer channel to exceed its critical time. From the perspective of the control unit, burst mode occurs when the time contributed by the control unit in a transfer sequence is more than 32 microseconds. (Refer to the *Enterprise System Architecture/390 System 360 and System 370 I/O Interface Channel to Control Unit OEMI*.)

If the device configuration guidelines are followed for byte multiplexer channels on a general purpose model CPC, deferred accesses are minimized and data transfer sequences exceeding 32 microseconds are acceptable when large burst sizes are specified.

Most class-2 and class-3 devices that can operate in burst mode should be attached to block multiplexer channels for better performance.

I/O operations control

ESA/390 and z/Architecture I/O operations are performed by executing a channel program that consists of one or more chained Channel Command Words (CCWs). Each CCW contains a command and other information that is used by both the channel and control unit in executing the I/O operation.

Channel commands are segmented into six basic categories with many variations based on control unit type. A channel program is initiated and controlled by executing one or more of the ESA/390 and z/Architecture I/O instructions described below. I/O interruptions may result during the execution of a channel program to notify the CP of progress or completion.

Channel commands

The six basic channel commands are:

Write Initiates the transfer of data from main storage to an I/O device.

Read Initiates the transfer of data from an I/O device to main storage.

Read Backward

Initiates the transfer of data from an I/O device to main storage, storing data bytes in reverse order.

Control

Specifies operations such as set tape density, rewind tape, advance paper in a printer, or sound an audible alarm.

Sense Requests information from a control unit. The information contains unusual conditions detected during the last I/O operation and detailed device status.

Transfer in Channel (TIC)

Specifies the location in main storage where the next CCW in the channel program is to be fetched. The TIC command provides branching between CCWs in noncontiguous storage areas. A TIC command cannot specify a CCW containing another TIC command.

ESA/390 and z/Architecture mode I/O instructions

In ESA/390 mode or z/Architecture mode, any CP can initiate I/O operations with any I/O device and can handle I/O interruptions from any I/O device. Each I/O device is assigned a unique device number, and is associated with one subchannel.

The CPs communicate with devices by specifying the appropriate subchannel. The subchannel uses the assigned device address to communicate with the device over one or more channel paths. The device number provides a path-independent means to refer to a device for use in operator messages or at the time of IPL.

For descriptions of these instructions, refer to the *Enterprise System Architecture/390 Principles of Operation* or *z/Architecture Principles of Operation* manual.

The I/O instructions for operation in ESA/390 mode or z/Architecture mode are:

- Start Subchannel (SSCH)
- Test Subchannel (TSCH)
- Clear Subchannel (CSCH)
- Halt Subchannel (HSCH)
- Resume Subchannel (RSCH)
- Store Subchannel (STSCH)
- Modify Subchannel (MSCH)
- Test Pending Interruption (TPI)
- Reset Channel Path (RCHP)
- Set Channel Monitor (SCHM)
- Store Channel Report Word (STCRW)
- Cancel I/O (XSCH).

The SSCH instruction specifies an operation request block, which designates the channel program.

Chaining operations

Following the transfer of information over a channel designated by a Channel Command Word (CCW), an operation initiated by the Start Subchannel (SSCH) instruction can be continued by fetching a new CCW. Fetching a new CCW immediately following the completion of the previous CCW is called **chaining**. Chaining is described in more detail in the *Enterprise System Architecture/390 Principles of Operation* or *z/Architecture Principles of Operation*.

CCWs located in contiguous areas of central storage (successive doubleword locations) can be chained. Chains of CCWs located in noncontiguous storage areas can be coupled for chaining purposes by using a Transfer in Channel command. All CCWs in a chain refer to the I/O device specified in the original instruction.

The type of chaining (data or command) is specified by chain-data and chain-command flag bits in the CCW.

Data chaining

When the data transfer specified by the current CCW is finished, data

chaining causes the operation to continue by fetching a new CCW and using the storage area defined by the new CCW. Execution of the operation at the I/O device is not affected.

Command chaining

Each time a new CCW is fetched during command chaining, a new I/O operation is specified. The new operation is initiated when the device end signal for the current operation is received, unless suspension is specified in the new CCW. When command chaining takes place, the completion of the current operation does not cause an I/O interruption.

I/O interruptions

I/O interruptions report the completion of I/O operations to the CPs, error and time-out conditions, and progress.

Ending status information about the operation is available to the control program at the end of the I/O operation. When an I/O operation is completed, an I/O interruption request is sent to a central processor. When the request is honored, an I/O interruption occurs and places the central processor under control of the I/O new program status word (PSW). Until an I/O interruption condition is honored, it is called a pending I/O interruption.

Errors detected by the channel subsystem are reported to the CPs as I/O interruptions or machine-check interruptions. I/O interruptions report the following hardware-related conditions:

- Interface Control Check (IFCC) - For example, interface tag errors and time-outs.
- Channel Control Check (CCC) - For example, parity, decode, or control errors.
- Channel Data Check (CDC) - For example, a parity error detected in central storage.

Machine-check interruptions include the following:

- Unrecoverable errors (retry is unsuccessful).
- Persistent errors (retry can be attempted, but the error threshold is exceeded).
- Serious channel element errors that require immediate reporting or cannot be reported as an IFCC or CCC with an I/O interruption.

Resets

An I/O system reset is issued to all channels, and the channels signal a system reset to all attached I/O devices.

An I/O system reset:

- Stops all subchannel operations.
- Resets interruptions and status in all subchannels.

An I/O system reset occurs as part of:

- Channel subsystem power-on reset.
- Initial program load.
- System reset.

A channel issues a selective reset to a specific I/O device in response to an IFCC, CCC, or as part of execution of the clear subchannel instruction. The status of the specific device is reset.

I/O interface protocol

The I/O interface protocol is determined by the interface sequencing operations selected for specific control units and their associated devices that are attached to the channel.

Channel-to-Channel connection

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 system, 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 I/O device.

The CTC is selected and responds in the same manner as any I/O device. It differs from other I/O devices in that it uses commands to open a path between the two channels it connects, and then synchronizes the operations performed between the two channels.

ESCON CTC support: The parallel I/O CTC architecture defines two operating modes for CTC communication: basic mode and extended mode. ESCON CTC support for both of these modes is available.

ESCON channels (using link-level and device-level protocols): You can achieve ESCON channel-to-channel connections between CPCs with ESCON or FICON Express channels if one of the ESCON channels is defined to operate in channel-to-channel (CTC) mode.

ESCON CTC type channels can talk with FICON Express FCV type channels through a FICON bridge only, not with FICON FC type channels.

ESCON channels that operate in CTC mode (extended mode or basic mode) can be defined as shared ESCON channels. For more information, refer to “Multiple Image Facility (MIF)” on page 40.

For detailed information about the ESCON channel-to-channel adapter, refer to *Enterprise Systems Architecture/390 ESCON Channel-to-Channel Adapter*.

Channel time-out functions

The optional time-out function described here applies only to ESCON channels that attach to a 9034 ESCON converter channel.

Each channel path has I/O interface time-out functions that time the control unit delays in completing the following I/O interface sequences:

- A 6-second time-out for all selection and status presentation sequences. A time-out occurs if the sequence is not complete within 6 seconds.
- A 30-second time-out for data transfer. A time-out occurs if a byte of data is not transferred within 30 seconds.

If a time-out occurs, the channel terminates the I/O request to the control unit and generates an IFCC interruption.

The time-out function detects malfunctions in control units and I/O devices that can cause the channel path to be unusable to other control units and I/O devices. The time-out function is specified as active or inactive for a device by IOCP when the IOCDS is created.

Control unit (CU) priority on an I/O interface

CU priority on an I/O interface applies only to ESCON channels attached to a 9034 ES connection converter channel.

CU priority on the I/O interface of a channel depends on the order in which they were attached. If the CUs are connected to the “select out” line, the first CU has the highest priority. If the CUs are attached to the “select in” line, the priority sequence is reversed. CUs attached to the “select out” line have priority over CUs attached to the “select in” line.

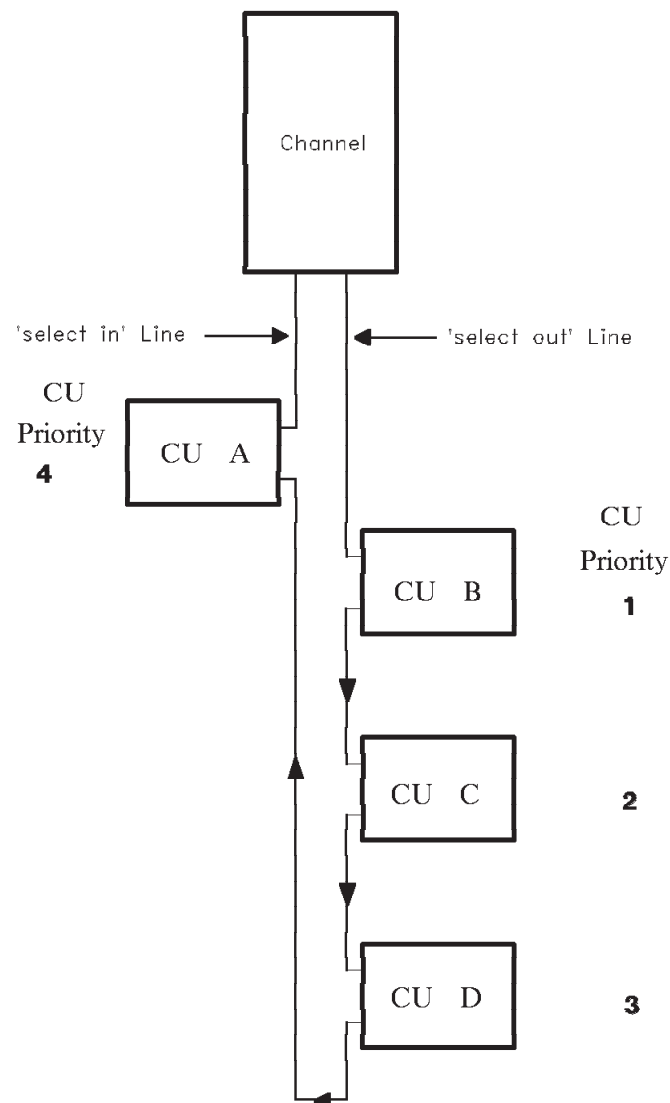


Figure 6. Control Unit (CU) priority on ESCON channels attached to a 9034 ES connection converter

Dynamic reconnection

The channel subsystem permits dynamic reconnection of I/O devices that have the dynamic-reconnection feature installed and that are set up to operate in a multipath mode, such as the IBM 3390 Direct Access Storage Model A14 or A22. Dynamic reconnection allows the device to reconnect and continue a chain of I/O operations using the first available channel path (one of as many as eight possible channel paths defined in an IOCP parameter). The selected path is not necessarily the one used initially in the I/O operation.

ESCON channel performance

Channel subsystem performance can be examined by observing two measurements:

- Response time (the amount of time taken to complete an I/O operation).
- Throughput (the number of I/O operations an I/O subsystem can complete in a given amount of time).

Channel subsystem response time and throughput can be divided into four major components:

- **Queuing and setup time**

- The time taken for a channel path, control unit, and device to become available.
- The time taken for a channel to send the I/O operation commands to the control unit.

- **Control unit and device time**

The time required by the control unit and device to prepare for the transfer of data for the I/O operation. For example, a non-cached DASD control unit may have to wait for the DASD's seek and latency times before being ready to accept or send data.

- **Data transfer time**

The time it takes to transfer the data for the I/O operation.

- **Completion time**

The time it takes for the channel and control unit to post the status of and end the I/O operation.

Factors affecting ESCON channel performance

Factors that affect the various components of performance include:

- Synchronous or nonsynchronous type of operation
- Data transfer rate
- Attached device characteristics
- Channel subsystem workload characteristics.

Synchronous and nonsynchronous I/O operation: For detailed information about concepts described in this section, refer to *Storage Subsystem Library Introduction to Nonsynchronous Direct Access Storage Subsystems*.

Synchronous operation

Most DASD devices in a parallel environment transfer data synchronously. Synchronous operation requires that the channel, control unit, and device be active at the same time.

All work involved in ending an operation and advancing to the next operation must be completed before the DASD head reaches the next record (commonly referred to as the inter-record gap). If this does not

occur, a rotational positional sensing/sensor (RPS) miss or an overrun is generated and the operation must wait for one DASD revolution before continuing.

Nonsynchronous operation

Nonsynchronous operation removes the requirements of synchronous operation. During nonsynchronous operation, the channel, control unit, and device do not have to be active at the same time to perform an I/O operation; thereby:

- Increasing DASD storage potential (by reducing inter-record gap).
- Allowing the channel and control units to be separated by longer distances.
- Eliminating command overruns.
- Reducing response time (by reducing RPS misses).
- Permitting the channel to perform other operations during the time it would normally wait for the device (this increases the throughput of the system).

Extended count key data (ECKD) channel programs are required to gain the benefits of nonsynchronous I/O operations. Count key data (CKD) channel programs are supported, but without the benefit of nonsynchronous operation. CKD channel-program performance could be degraded relative to ECKD channel programs in a nonsynchronous environment.

Data transfer rate: One of the factors that affects channel performance is the data transfer rate. The I/O subsystem data rate is the data transfer rate between processor storage and the device during an I/O operation.

The I/O subsystem data rate is made up of three components:

- **Channel data rate**

The rate that the channel transfers data between the transmission link and processor storage during an I/O operation. For ESCON channels, the link speed is 20 MBps and the channel data rate is 17 MBps at 0 distance. The data rate increases with distance.

- **Control unit data rate**

The rate that the control unit transfers data between the control unit and the transmission link during an I/O operation.

- **Device data rate**

The rate of data transfer between the control unit and the device. This rate depends on the control unit and device you use.

The **I/O subsystem data rate** is the lowest of the channel data rate, the control unit data rate, and the device data rate. In cases where the data comes from the control unit or is stored on the control unit and not directly to the device (for example, a cache read), the I/O subsystem data rate is the lower of the two: channel data rate or the control unit data rate.

The I/O subsystem data rate affects only the data transfer portion of the response time for an I/O operation. Response time and throughput both improve (response time decreases and throughput increases).

I/O device characteristics: The characteristics of devices attached to a channel subsystem can have a substantial effect on performance. Device characteristics such as caches, buffers, and data transfer rates all affect response time and throughput.

Channel subsystem workload characteristics: The performance of a specific I/O configuration varies based on the workload characteristics of that configuration. Two significant factors that determine workload characteristics and affect response time and throughput are channel program characteristics and cache-hit rates.

Channel program characteristics

Channel program characteristics affect channel subsystem performance. ESCON channel subsystems using link-level and device-level protocols perform nonsynchronous data transfers, and should use extended count key data (ECKD) channel programs.

Count key data (CKD) channel programs run in an ESCON environment, but may increase response times and reduce throughput due to lost DASD rotations.

Channel programs that contain indirect data address words (IDAWs), Transfer in Channel commands (TICs), and chained data commands, or that have poorly-aligned data boundaries, cause longer storage-response and increase channel subsystem response times.

Chained data commands increase response time due to an additional interlocked exchange between the channel and control unit. Refer to “ESCON performance characteristics” for more information.

The amount of data to be transferred per I/O operation affects throughput. As the amount of data transferred per I/O operation increases (the ratio of data transferred to overhead improves), throughput improves.

Cache-hit rates

For control units which implement caches, cache-hit rates affect the channel subsystem performance. As the cache-hit rate increases, response time and throughput improve. The cache-hit rate is the percentage of times when data needed for a read operation is in the control unit's cache. For example, a cache-hit rate of 70% means that the required data is in the cache for 7 out of 10 read operations.

The cache-hit rate is significant because data is transferred out of the cache at the control unit's maximum data transfer rate, while data from the device is transferred at lower device speeds. This means that the higher the cache-hit rate, the better the response time and the better the throughput.

ESCON performance characteristics

This section describes the performance characteristics of ESCON channels and ESCON channels attached to a 9034.

With ESCON channels you need to consider the distance between the channel and the control unit since this affects the setup and completion times of an I/O operation. As the distance between the channel and the control unit increases, the response time increases and the throughput decreases. Channel and control unit utilization also increases as distance between the channel and control unit increases.

The speed of data transfer through fiber optic cable is subject to the Propagation delay time is determined by two factors: the speed of light through the optical fiber

(which is fixed), and the length of the fiber optic link. Propagation delay time increases as the distance between elements in a fiber optic environment increase.

Interlocked exchange affects response time. Interlocked exchange requires that the channel (or control unit) wait for a response from the control unit (or channel) before proceeding with the next step of an I/O operation. As distance increases, the interlocked-exchange response time increases because of longer propagation delay times.

The throughput and response time for a shared ESCON channel are comparable to that of an unshared ESCON channel with comparable workload.

OSA channels

OSA channels include all OSA-Express2 and OSA-Express3 features.

Note: Unless noted differently, throughout this section, the term “OSA features” refers to all the OSA-Express3 and OSA-Express2 features.

Supported CHPID types

OSA channels support the following modes of operation:

- CHPID type OSD
 - OSA-Express2 or OSA-Express3 features is running in QDIO mode.
QDIO mode is the preferred architecture on System z9 and System z10 for high-speed communication, helping to reduce host interruptions and improve response time.
 - TCP/IP traffic when Layer 3
 - Protocol-independent when Layer 2
- CHPID type OSE
 - OSA-Express2 or OSA-Express3 features is running in non-QDIO mode.
 - SNA/APPN/HPF and/or TCP/IP passthru (LCS)
- CHPID type OSC
 - OSA-Integrated Console Controller (OSA-ICC)
 - TN3270E, non-SNA DFT to IPL CPCs and LPARs
 - Operating system console operations
- CHPID type OSN
 - OSA-Express for NCP
 - Supports channel data link control (CDLC) protocol. This provides connectivity between System z operating systems and IBM Communication Controller for Linux (CCL).
CCL allows you to keep data and applications on the mainframe operating systems while moving NCP function to Linux on System z. CCL on System z helps to improve network availability by replacing token-ring networks and ESCON channels with an Ethernet network and integrated LAN adapters on System z10, OSA-Express3 GbE or 1000BASE-T Ethernet features, or OSA-Express2 GbE or 1000BASE-T Ethernet features.
 - Requires the configuring to be done on a port-by-port basis
 - Used exclusively for internal communication, LPAR-to-LPAR
 - Supports concurrent LIC updates
 - Supported in the z/OS, z/VM, z/VSE, TPF, z/TPF, and Linux on System z environments.

For more detailed information on these operating modes, refer to *System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*.

OSA/SF

The Open Systems Adapter Support Facility (OSA/SF) is a host-based tool to support and manage the OSA features operating in QDIO (CHPID type OSD), non-QDIO mode (CHPID type OSE), or for OSA-Express for NCP (CHPID type OSN). The OSA/SF is used primarily to manage all OSA ports, configure all OSA non-QDIO ports, and configure local MACs.

One OSA/SF application can communicate with all OSA features in a hardware complex. OSA/SF communicates with an OSA feature through a device predefined on the OSA feature. The device type is OSAD.

OSA/SF is not required to set up the OSA features in QDIO mode (CHPID type OSD). However, it can be used to set up MAC addresses and set adapter speed. For channels (CHPID type OSN), OSA/SF does not provide any configuration management capabilities but provides capabilities only for operations management.

OSA/SF includes a Java-based Graphical User Interface (GUI) in support of the client application. The Java GUI is independent of any operating system/server (transparent to operating system), and is expected to operate wherever the Java 1.4 runtimes are available.

Interoperability testing has been performed for Windows 2000, Windows XP, and Linux on System z.

Use of the GUI is optional; a REXX command interface is also included with OSA/SF. OSA/SF has been, and continues to be, integrated in z/OS, z/VM, and z/VSE and runs as a host application. For OSA/SF, Java GUI communication is supported via TCP/IP only.

The Layer 3 OSA Address Table (OAT) displays all IP addresses registered to an OSA port.

OSA/SF has the capability of supporting virtual Medium Access Control (MAC) and Virtual Local Area Network (VLAN) identifications (IDs) associated with OSA-Express2 and OSA-Express3 features configured as a Layer 2 interface.

These OSA/SF enhancements are applicable to CHPID type OSD, OSE, and OSN.

For more detailed information on OSA/SF, refer to *System z10, System z9 and eServer zSeries Open Systems Adapter-Express Customer's Guide and Reference*.

OSA-Express3 features

The OSA-Express3 features are designed for use in high-speed enterprise backbones, for local area network connectivity between campuses, to connect server farms to z10 BC, and to consolidate file servers onto z10 BC. The workload can be Internet Protocol (IP) based or non-IP based. All OSA-Express3 features are hot-pluggable.

OSA-Express3 provides the following enhancements compared to OSA-Express2:

- Double the port density on a single feature. Although the maximum number of combined OSA features supported on the z10 BC server is 24, when all ports are

OSA-Express3, a maximum of 96 OSA-Express3 ports of LAN connectivity are allowed compared to a maximum of 48 when all ports are OSA-Express2. With more ports per feature, the number of I/O slots and I/O resources can be reduced.

- Optimized latency mode helps to improve the performance of z/OS workloads by minimizing response times for inbound and outbound data when servicing remote clients. This only applies to OSA-Express3 and CHPID type OSD (QDIO).
- Reduced latency and improved throughput for standard frames (1492 bytes) and jumbo frames (8992 bytes) due in part by the PCIe adapter and the hardware data router.

OSA-Express3 features have PCIe adapters. The PCIe adapter supports 8 lanes and transfers data at a rate of 2.5 Gbps per lane (data transfer rate of 2 Gbps in each direction). There are two ports on each PCIe adapter.

Note: For operating systems to exploit all ports on the two port or four port features, software updates must be installed.

The hardware data router provides a direct host memory to LAN flow without firmware intervention.

The OSA-Express3 features includes:

- **OSA-Express3 Gigabit Ethernet (GbE) LX (FC 3362)**

OSA-Express3 GbE LX has four ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having two ports. Supports CHPID types: OSD and OSN

The OSA-Express3 GbE LX uses a 9 micron single mode fiber optic cable with an LC duplex connector and a link data rate of 1000 Mbps (1 Gbps). However, OSA-Express3 GbE LX also accommodates the reuse of existing multimode fiber (50 or 62.5 micron) when used with a pair of mode conditioning patch (MCP) cables. It is designed to support unrepeated distances of up to 5 km (3.1 miles). If using MCP cables, the supported unrepeated distance is 550 meters (1804 feet).

- **OSA-Express3 Gigabit Ethernet (GbE) SX (FC 3363)**

OSA-Express3 GbE SX has four ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having two ports. Supports CHPID types: OSD and OSN

The OSA-Express3 GbE SX uses a 50 or 62.5 micron multimode fiber optic cable with an LC duplex connector and a link data rate of 1000 Mbps (1 Gbps). The supported unrepeated distances vary:

- With 50 micron fiber at 500 MHz-km: 550 meters (1804 feet)
- With 62.5 micron fiber at 200 MHz-km: 273 meters (902 feet)
- With 62.5 micron fiber at 160 MHz-km: 220 meters (722 feet)

- **OSA-Express3 1000BASE-T Ethernet (FC 3367)**

OSA-Express3 1000BASE-T Ethernet has four ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having two ports. Supports CHPID types: OSD, OSE, OSC, and OSN

The OSA-Express3 1000BASE-T Ethernet uses a EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cable with an RJ-45 connector and a maximum length of 100 meters (328 feet). It supports a link data rate of 10, 100, or 1000 Mbps; half duplex and full duplex operation modes; and autonegotiations to other speeds.

- **OSA-Express3-2P 1000BASE-T Ethernet (FC 3369)**

OSA-Express3-2P 1000BASE-T Ethernet has two ports per feature occupying a single I/O slot. Each feature has one CHPID, with this CHPID having two ports. Supports CHPID types: OSD, OSE, OSC, and OSN

The OSA-Express3-2P 1000BASE-T Ethernet uses a EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cable with an RJ-45 connector and a maximum length of 100 meters (328 feet). It supports a link data rate of 10, 100, or 1000 Mbps; half duplex and full duplex operation modes; and autonegotiations to other speeds.

- **OSA-Express3 10 Gigabit Ethernet (GbE) Long Reach (LR) (FC 3370)**

OSA-Express3 10 GbE LR has two ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having one port. Supports CHPID types: OSD

OSA-Express3 10 GbE LR uses a 9 micron single mode fiber optic cable with an LC duplex connector and a link data rate of 10 Gbps. It is designed to support unrepeated distances of up to 10 km (6.2 miles).

OSA-Express3 10 GbE LR does not support auto-negotiation to any other speed. It supports 64B/66B coding.

- **OSA-Express3 10 Gigabit Ethernet (GbE) Short Reach (SR) (FC 3371)**

OSA-Express3 10 GbE SR has two ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having one port. Supports CHPID types: OSD

OSA-Express3 10 Gigabit Ethernet SR uses a 50 or 62.5 micron multimode fiber optic cable with an LC duplex connector and a link data rate of 10 Gbps. The supported unrepeated distances vary:

- With 50 micron fiber at 2000 MHz-km: 300 meters (984 feet)
- With 50 micron fiber at 500 MHz-km: 82 meters (269 feet)
- With 62.5 micron fiber at 200 MHz-km: 33 meters (108 feet)

- **OSA-Express3-2P Gigabit Ethernet (GbE) SX (FC 3373)**

OSA-Express3-2P GbE SX has two ports per feature occupying a single I/O slot. Each feature has one CHPID, with this CHPID having two port. Supports CHPID types: OSD and OSN

The OSA-Express3-2P GbE SX uses a 50 or 62.5 micron multimode fiber optic cable with an LC duplex connector and a link data rate of 1000 Mbps (1 Gbps). The supported unrepeated distances vary:

- With 50 micron fiber at 500 MHz-km: 550 meters (1804 feet)
- With 62.5 micron fiber at 200 MHz-km: 273 meters (902 feet)
- With 62.5 micron fiber at 160 MHz-km: 220 meters (722 feet)

All OSA-Express3 features support full duplex operation and standard frames (1492 bytes) and jumbo frames (8992 bytes).

OSA-Express2 features

OSA-Express2 LAN adapters help to ensure you have a balanced system to satisfy the bandwidth demands of your applications. The OSA-Express2 features are hot-pluggable, support the Multiple Image Facility (MIF) that shares channels across logical partitions (LPARs), and can be defined as a spanned channel to be shared among LPARs within and across LCSSs. The maximum combined number of OSA features supported on the z10 BC server is 24.

The OSA-Express2 features can be carried forward to System z10 BC or they can be ordered.

OSA-Express2 features include:

- **OSA-Express2 Gigabit Ethernet (GbE) LX (FC 3364)**

OSA-Express2 GbE LX has two ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having one port. Supports CHPID types: OSD and OSN

The OSA-Express2 GbE LX uses a 9 micron single mode fiber optic cable with an LC duplex connector and a link data rate of 1000 Mbps (1 Gbps). However, OSA-Express2 GbE LX also accommodates the reuse of existing multimode fiber (50 or 62.5 micron) when used with a pair of mode conditioning patch (MCP) cables. It is designed to support unrepeated distances of up to 5 km (3.1 miles). If using MCP cables, the supported unrepeated distance is 550 meters (1804 feet).

- **OSA-Express2 Gigabit Ethernet (GbE) SX (FC 3365)**

OSA-Express2 GbE SX has two ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having one port. Supports CHPID types: OSD and OSN

The OSA-Express2 GbE SX uses a 50 or 62.5 micron multimode fiber optic cable with an LC duplex connector and a link data rate of 1000 Mbps (1 Gbps). The supported unrepeated distances vary:

- With 50 micron fiber at 500 MHz-km: 550 meters (1804 feet)
- With 62.5 micron fiber at 200 MHz-km: 273 meters (902 feet)
- With 62.5 micron fiber at 160 MHz-km: 220 meters (722 feet)

- **OSA-Express2 1000BASE-T Ethernet (FC 3366)**

OSA-Express2 1000BASE-T Ethernet has two ports per feature occupying a single I/O slot. Each feature has two CHPIDs, with each CHPID having one port. Supports CHPID types: OSD, OSE, OSC, and OSN

The OSA-Express2 1000BASE-T Ethernet uses a EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cable with an RJ-45 connector and a maximum length of 100 meters (328 feet). It supports a link data rate of 10, 100, or 1000 Mbps over a copper infrastructure; half duplex and full duplex operation modes; and autonegotiations to other speeds.

- **OSA-Express2 10 Gigabit Ethernet (GbE) Long Reach (LR) (FC 3368)**

OSA-Express2 10 GbE LR is used as the enterprise backbone, between campuses, to connect server farms to, and consolidate file servers on, your z10 BC server.

OSA-Express2 10 GbE LR has one CHPID per feature, with one port on that CHPID. Supports CHPID types: OSD

OSA-Express2 10 GbE LR uses a 9 micron single mode fiber optic cable with an SC duplex connector and a link data rate of 10 Gbps. It is designed to support unrepeated distances of up to 10 km (6.2 miles).

All OSA-Express2 features support full duplex operation and standard frames (1492 bytes) and jumbo frames (8992 bytes).

OSA-Express3 and OSA-Express2 supported functions

The following is a list of functions supported by the OSA-Express3 and OSA-Express2 features in QDIO mode (CHPID type OSD):

Note: Throughout this section, the term “OSA” refers to both OSA-Express3 and OSA-Express2.

Dynamic LAN idle

The OSA LAN idle timer value defines how long OSA will hold packets before presenting the packets to the host. The LAN idle function now allows the host OS to

dynamically update the existing LAN idle timer values (defined within the QIB) while the specific QDIO data device is in the QDIO active state.

Dynamic LAN idle timer function is exploited by z/OS V1.8 with PTFs or later and z/VM 5.3 or later for guest exploitation.

Network Traffic Analyzer Trace facility

The OSA-Express Network Traffic Analyzer Trace facility is a diagnostic tool used to copy frames as they enter or leave an OSA adapter for an attached host. This facility is controlled and formatted by the z/OS Communications Server, but the data is collected in the OSA at the network port. Because the data is collected at the Ethernet frame level, you can trace the MAC headers for packets. You can also trace ARP packets, SNA packets, and packets being sent to and from other users sharing the OSA, including other TCP/IP stacks, Linux on System z users, and z/VM users.

To enable the OSA-Express Network Traffic Analyzer Trace facility, you must be running with a minimum of z/OS V1.8 with PTFs or later.

Queued Direct I/O Diagnostic Synchronization (QDIOSYNC)

Queued Direct I/O Diagnostic Synchronization provides the ability to coordinate and simultaneously capture software (z/OS) and hardware (OSA) traces. This function allows the host operating system to signal the OSA feature to stop traces and allows the operator to capture both the hardware and software traces at the same time. You can specify an optional filter that alters what type of diagnostic data is collected by the OSA adapter. This filtering reduces the overall amount of diagnostic data collected and therefore decreases the likelihood that pertinent data is lost.

To use the Queued Direct I/O Diagnostic Synchronization facility, you must be running with a minimum of z/OS V1.8 with PTFs or later.

Dynamic link aggregation for the z/VM environment

This function dedicates an OSA port to the z/VM 5.3 or later operating system for link aggregation under z/VM Virtual Switch-controlled link aggregation. Link aggregation (trunking) is designed to allow you to combine multiple physical OSA ports of the same type into a single logical link. You can have up to eight OSA ports in one virtual switch. This increases bandwidth and permits nondisruptive failover in the event that a port becomes unavailable. This function also supports dynamic add/remove of OSA ports and full-duplex mode (send and receive).

This function is supported on all OSA-Express2 and OSA-Express3 features in Layer 2 mode in QDIO mode (CHPID type OSD).

Multiple Image Facility (MIF) and spanned channels

OSA features support the Multiple Image Facility (MIF) for sharing channels across LPARs. Then can be defined as a spanned channel to be shared among LPARs within and across LCSS.

QDIO data connection isolation

QDIO data connection isolation provides protection for workloads (servers and clients) hosted in a virtual environment from intrusion or exposure of data and processes from other workloads.

This function applies to the OSA-Express3 and OSA-Express2 features (CHPID type OSD) on System z10 and to OSA-Express2 features (CHPID type OSD) on System z9.

Layer 2 (Link Layer) support

OSA features can support two transport modes when using CHPID type OSD (QDIO): Layer 2 (Link Layer) and Layer 3 (Network or IP Layer). Layer 2 support can help facilitate server consolidation and will allow applications that do not use IP protocols to run on the z10 BC server.

640 TCP/IP stacks

Increasing the TCP/IP stacks allows you to host more Linux on System z images. OSA supports 640 TCP/IP stacks or connections per dedicated CHPID, or 640 total stacks across multiple LPARs using a shared or spanned CHPID when priority specification is disabled.

Large send

Large send improves performance by offloading TCP packet processing from the host to the TCP/IP stack. Offloading allows the host to send IP datagrams up to 60K in size. The IP datagram is controlled by the host TCP/IP stack. Sending larger data blocks reduces host processor utilization while increasing network efficiencies.

Concurrent LIC update

Allows you to apply LIC updates without requiring a configuration off/on, thereby minimizing the disruption of network traffic during the update. This applies to CHPID types OSD and OSN.

Layer 3 virtual MAC

The z/OS Layer 3 Virtual MAC (VMAC) function simplifies the network infrastructure and facilitates IP load balancing when multiple TCP/IP instances are sharing the same OSA port or Media Access Control (MAC) address. With Layer 3 VMAC support, each TCP/IP instance has its own unique "virtual" MAC address instead of sharing the same universal or "burned in" OSA MAC address. Defining a Layer 3 VMAC provides a way for the device to determine which stack, if any, should receive a packet, including those received for IP addresses that are not registered by any TCP/IP stack. With Layer 3 VMAC in a routed network, OSA appears as a dedicated device to the particular TCP/IP stack, which helps solve many port-sharing issues.

Layer 3 Virtual MAC function is supported by z/OS V1.8 with PTFs or later and z/VM V5.3 or later for guest exploitation.

Jumbo frames

When operating at 1 Gbps (fiber or copper) and 10 Gbps (fiber), use of jumbo frames (8992 bytes) are supported.

HiperSockets

HiperSockets "network within the box" functionality allows high speed any-to-any connectivity among OS images within the z10 BC server without requiring any physical cabling. This "network within the box" concept minimizes network latency and maximizes bandwidth capabilities between z/VM, Linux on System z, z/VSE, and z/OS images (or combinations of these) to enable optimized e-business and ERP solutions within a single server. These images can be first level (i.e. directly under LPAR), or second level images (i.e. under z/VM). Up to 16 separate internal LANs can be configured within a server thereby allowing OS images to be grouped according to the function they provide. These groupings are independent of sysplex affiliation.

Separate HiperSocket LANs are mainly required if some logical partitions need to be isolated from other logical partitions. Each LAN is configured as an CHPID type IQD.

In addition the number of communication queues is 4096 and each queue can have three subchannels. If you want the internal LANs shared between partitions in different LCSSs then the channel must be spanned. For more information on spanned channels, refer to “Spanned channels” on page 41.

Broadcast support

Internet Protocol Version 6 (IPv6) broadcast packets are supported over HiperSockets internal LANs. TCP/IP applications that support IPv6 broadcast, such as OMROUTE when running Routing Information Protocol Version 1 (RIPv1), can send and receive broadcast packets over HiperSockets interfaces. IPv4 and IPv6 broadcast support is available for HiperSockets on z/OS, z/VM V5.3 or later, and Linux on System z. Refer to <http://www.ibm.com/developerworks/linux/linux390/> for more information on Linux on System z support.

IPv6 support

HiperSockets supports Internet Protocol Version 6 (IPv6). IPv6 expands the IP address space from 32 bits to 128 bits to enable a greater number of unique IP addresses in support of the proliferation of devices, such as cell phones and PDAs, now connecting to the Internet.

IPv4 and IPv6 support is available for HiperSockets on z/OS, z/VM, and Linux on System z.

VLAN support

Virtual Local Area Networks (VLANs), IEEE standard 802.1q, is supported in HiperSockets in a Linux on System z environment. VLANs increase bandwidth and reduce overhead by allowing networks to be organized for more optimum traffic flow. The network is organized by traffic patterns rather than physical location. This allows traffic to flow on a VLAN connection over HiperSockets and between HiperSockets and OSA.

HiperSockets Network Concentrator

HiperSockets Network Concentrator simplifies network addressing between HiperSockets and OSA allowing seamless integration of HiperSockets-connected operating systems into external networks, without requiring intervening network routing overhead, thus helping to increase performance and simplify configuration.

HiperSockets Network Concentrator is implemented between HiperSockets, OSA, and Linux on System z. The Network Concentrator provides support for unicast, broadcast, and multicast. For more information, refer to <http://www.ibm.com/developerworks/linux/linux390/>.

HiperSockets Network Traffic Analyzer

The HiperSockets Network Traffic Analyzer Trace facility is used to diagnose problems in a HiperSockets network. As data flows over an IQD channel, the HiperSockets Network Traffic Analyzer captures and analyzes each packet. The captured data can be displayed immediately or written to a file.

The captured data includes packets being sent to and from other users sharing the HiperSockets channel, such as logical partitions with z/OS, Linux on System z, z/VSE, or z/VM and z/VM guests.

To use this function, the level of authorization for the HiperSockets network traffic analyzer must be selected. This authorization determines the scope of the tracing. Then a HiperSocket tracing device must be activated on your system. This is performed by the operating system of the owning partition.

Setting the authorization level is performed on the Support Element using the Network Traffic Analyzer Authorization task. The levels of authorization are as follows:

- No traffic on any IQD channel for the selected server can be traced
- No traffic on the selected IQD channel can be traced
- All traffic on the selected IQD channel can be traced. (This traces all traffic flowing between all the logical partitions using this IQD CHPID.)
- Customized traffic flow between selected logical partitions can be traced.

From the Customize a HiperSockets NTA Logical Partition Authorization List window, select the logical partition that will be authorized to set up, trace, and capture the HiperSockets network traffic. Then select all eligible partitions to be traced. Only the traffic flowing between the selected eligible partition or partitions will be traced.

The Support Element issues security logs to create an audit trail of the HiperSockets network traffic analyzer tracing activity.

Layer 2 (Link Layer) support

HiperSockets supports two transport modes on the z10 EC: Layer 2 (Link Layer) and Layer 3 (Network and IP Layer). HiperSockets in Layer 2 mode can be used by Internet Protocol (IP) Version 4 or Version 6 and non-IP protocols (such as AppleTalk, DECnet, IPCX, NetBIOS, or SNA).

Each HiperSockets device has its own Layer 2 MAC address and allows the use of applications that depend on a Layer 2 address such as DHCP servers and firewalls. LAN administrators can configure and maintain the mainframe environment in the same fashion as they do in other environments. This eases server consolidation and simplifies network configuration.

The HiperSockets device performs automatic MAC address generation to create uniqueness within and across logical partitions and servers. MAC addresses can be locally administered, and the use of Group MAC addresses for multicast and broadcasts to all other Layer 2 devices on the same HiperSockets network is supported. Datagrams are only delivered between HiperSockets devices using the same transport mode (Layer 2 with Layer 2 and Layer 3 with Layer 3).

A HiperSockets Layer 2 device may filter inbound datagrams by VLAN identification, the Ethernet destination MAC address, or both. This reduces the amount of inbound traffic, leading to lower CPU utilization by the operating system.

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

HiperSockets Layer 2 support is exclusive to System z10 and is supported by Linux on System z and by z/VM guest exploitation.

For hardware and software requirements, refer to the z/OS, z/VM, z/VSE subsets of the 2098DEVICE Preventive Service Planning (PSP) bucket prior to installing System z10 BC.

Multiple Write facility

HiperSockets allows the streaming of bulk data over a HiperSockets link between LPARs. The receiving LPAR can process a much larger amount of data per I/O interrupt. This function is transparent to the operating system in the receiving LPAR. HiperSockets Multiple Write facility, with fewer I/O interrupts, is designed to reduce CPU utilization of the sending and receiving LPAR.

HiperSockets Multiple Write facility is supported in the z/OS environment.

Chapter 6. Sysplex functions

This chapter describes the following z10 BC sysplex functions:

- “Parallel Sysplex”
- “Coupling facility” on page 80
- “System-managed CF structure duplexing” on page 86
- “Geographically Dispersed Parallel Sysplex (GDPS)” on page 87
- “Intelligent Resource Director (IRD)” on page 88.

Parallel Sysplex

IBM Parallel Sysplex makes use of a broad range of hardware and software products to process, in parallel, a transaction processing workload across multiple z/OS images with direct read/write access to sharing data.

The Parallel Sysplex allows you to manage a transaction processing workload, balanced across multiple z/OS images running on multiple Central Processor Complexes (CPCs), as a single data management system. It also offers workload availability and workload growth advantages.

The Parallel Sysplex enhances the capability to continue workload processing across scheduled and unscheduled outages of individual CPCs participating in a Parallel Sysplex using a coupling facility by making it possible to dynamically reapportion the workload across the remaining active Parallel Sysplex participants. Additionally, you can dynamically add processing capacity (CPCs or LPs) during peak processing without disrupting ongoing workload processing.

z10 BC CPC support for the Parallel Sysplex consists of having the capability to do any or all of the following:

- Configure IC links and define them as CHPID type ICP (peer link - connects to another IC)
- Install ISC-3 links and define them as type CFP (peer link - connects to another ISC-3)
- Install ICB-4 (connects System z10 to System z9, z990, or z890) links and define them as CHPID type CBP (peer link - connects to another ICB-4)
- Install 12x IFB links (connects System z10 to System z10 or System z10 to a System z9) and define them as CHPID type CIB (connects to another 12x IFB).
- Install 1x IFB links (connects System z10 to System z10) and define them as CHPID type CIB (connects to another 1x IFB).
- Define, as an LPAR, a portion or all of the CPC hardware resources (CPs, ICFs, storage, and coupling connections) for use as a coupling facility that connects to z/OS or another CF.
- Connect to a coupling facility for data sharing or resource sharing.

The z10 BC models provide the following support for the Parallel Sysplex:

- The z10 BC's Parallel Sysplex support consists of supporting coupling facilities on z10 BC, supporting attachment to remote coupling facilities via various type of coupling links, supporting Server Time Protocol (STP) or attachment to a Sysplex Timer for purposes of sysplex-consistent time, and supporting various ancillary CPC functions used by Parallel Sysplex support.

Note: If you want to implement a Parallel Sysplex between two sites, the fiber distance has been increased to 100 km. Refer to page 88 for more information.

- Internal Coupling links can be used to connect either z/OS images to coupling facilities (CFs) or CF images to other CF images within a z10 BC CPC. IC links have the advantage of providing CF communication at memory speed and do not require physical links.

These various interconnect formats provide the connectivity for data sharing between a coupling facility and the CPCs or logical partitions (LPs) directly attached to it.

Parallel Sysplex coupling link connectivity

z10 BC supports IC, ICB-4, ISC-3 and IFB for passing information back and forth over high speed links in a Parallel Sysplex environment. These technologies are all members of the family of coupling connectivity options available on z10 BC. With Server Time Protocol (STP), coupling links can also be used to exchange timing information. Refer to “Server Time Protocol (STP)” on page 83 for more information about Server Time Protocol. Refer to Table 11 for a summary of the coupling link options.

Table 11. Coupling link options

Link Type	Name	Communication Use	Maximum Links ¹
IFB ^{2, 3}	InfiniBand fiber connection between coupling facility and operating system	z10 EC, z10 BC z9 EC, z9 BC	12
ICB-4 ³	Copper connection between operating system and coupling facility	z10 EC, z10 BC z9 EC, z9 BC, z890, z990	12
ISC-3 ⁴	Fiber connection between operating system and coupling facility	z10 EC, z10 BC z9 EC, z9 BC, zSeries	48
IC	Internal coupling channel	Internal communication between CFs and z/OS LPARs	32

Note:

1. The maximum number of coupling links combined (ICs, ICB-4s, IFBs, and active ISC-3 links) cannot exceed 64 per server.
2. Each IFB link supports up to 16 CHPIDs (CHPID type CIB) per fanout. z10 BC supports up to 12 IFB links.
3. Maximum number of ICB-4 and IFB coupling links combined cannot exceed 12 per server.
4. ISC-3 peer mode only.

Notes:

1. ISC-3, ICB-4, and IFB links require a point-to-point connection (direct channel attach between a CPC and a coupling facility).
2. ISC-3, ICB-4, and IFB links can be redundantly configured (two or more ISC-3, ICB-4, or IFB links from each CPC to enhance availability and avoid extended recovery time).

Peer mode

Peer mode applies to ICBs (ICB-4 to ICB-4), IFBs (IFB to IFB), and ISCs (ISC-3 to ISC-3) operate only in peer mode. Peer mode provides both sender and receiver capability on the same link. Peer links provide up to seven expanded buffer sets.

Refer to Figure 7 for an illustration of these coupling links.

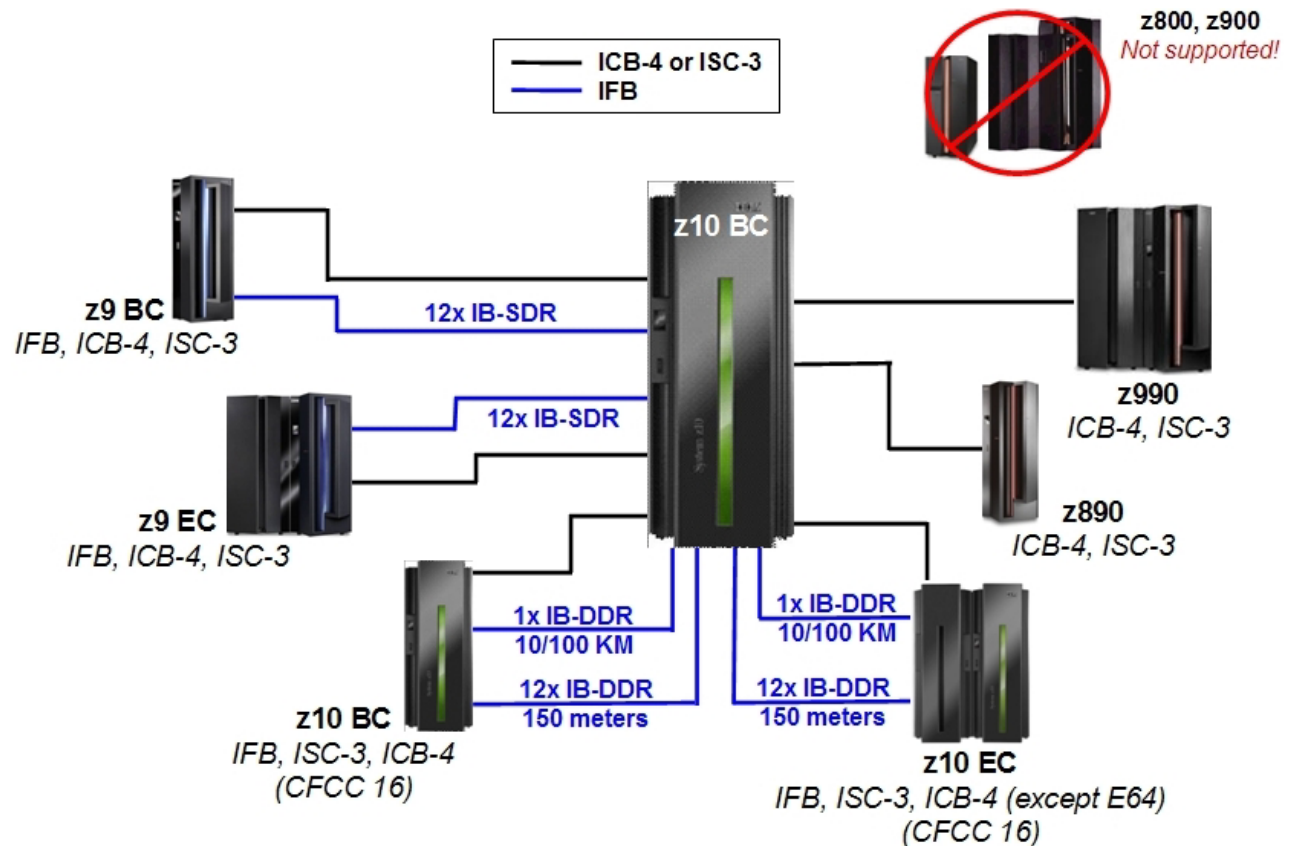


Figure 7. Coupling link connectivity

When coupling within a z10 BC server, the IC channel can be shared among several LPARs and one coupling facility partition.

ISC-3 links

The ISC-3 feature, with a link data rate of 2 Gbps, is a member of the family of coupling link options available on z10 BC. The ISC-3 feature is used by coupled systems to pass information back and forth over high speed links in a Parallel Sysplex environment. When STP is enabled, ISC-3 links can be used to transmit STP timekeeping information and can also be defined as timing-only links to other z10 BC servers, as well as z10 EC, z9 EC, z9 BC, z900, z890, and z990 servers.

The z10 BC ISC-3 feature is compatible with ISC-3 features on z10 EC, z9 EC, z9 BC, z900, z890, and z990 servers. ISC-3 (CHPID type CFP) can be defined as a spanned channel and can be shared among LPARS within and across LCSSs. z10 BC supports 48 ISC-3 links in peer mode, 12 features (four links per feature).

The ISC-3 feature is composed of:

- One Mother card (ISC-M), FC 0217

- Two Daughter cards (ISC-D), FC 0218.

Each daughter card has two ports or links, for a total of four links per feature. Each link is activated by using the Licensed Internal Code Configuration Control (LICCC) and can only be ordered in increments of one. The ISC-D is not orderable. Extra ISC-M cards can be ordered, up to a maximum of 12 or the number of ISC-D cards, whichever is less. When the quantity of ISC links (FC 0219) is selected, the appropriate number of ISC-M and ISC-D cards is selected by the configuration tool. Each port operates at 2 Gbps.

Each port utilizes a Long Wavelength (LX) laser as the optical transceiver, and supports use of a 9/125-micrometer single mode fiber optic cable terminated with an industry standard small form factor LC duplex connector. The ISC-3 feature accommodates reuse (at reduced distances) of 50/125-micrometer multimode fiber optic cables when the link data rate does not exceed 1 Gbps. A pair of mode conditioning patch cables are required, one for each end of the link.

ICB-4 links

The Integrated Cluster Bus (ICB-4) feature (FC 3393) provides a connection between two z10 BC servers or between an z10 BC server and a z10 EC, z9 EC, z9 BC, z990, or z890 server. ICB-4s can be defined as a spanned channel and can be shared among LPARs within and across LCSSs. An ICB-4 connection consists of one link that attaches directly to an STI port in the system, does not require connectivity to a card in the I/O drawer, and provides one output port to support ICB-4 to ICB-4 connectivity. One ICB-4 connection is required for each end of the link.

ICB-4 supports a link data rate of 2 GigaBytes per second (GBps) and carries traffic over 10 meters (33 feet) of copper cable, of which three meters are reserved for intraserver connection. Only one cable is required for each pair of ICB features. z10 BC requires new ICB-4 cables.

ICB-4 channel paths are defined as CHPID type CBP. ICB-4 channel paths allow a single channel to perform both send and receive functions on the same channel.

When STP is enabled, ICB-4 links can be used to transmit STP timekeeping information and can also be defined as timing-only links to other z10 BC, z10 EC, z9 EC, z9 BC, z990, and z890 servers. Refer to Figure 5 on page 37 for an overview of the ICB link connections.

ICB cables are ordered to match the quantity of ICBs on order. The quantity of ICB cables can be reduced, but cannot exceed the quantity of ICB functions on order. An ICB-4 cable connecting a z10 BC to a System z9 is different from an ICB-4 cable connecting a z10 BC to a z10 BC or z10 EC. ICB-4 cables are provided as a features on z10 BC.

Note: ICB-4 cables are not supported in a non-raised floor environment.

InfiniBand (IFB) links

There are four types of InfiniBand optical links supported by z10 BC, each supporting a point-to-point topology:

- 12x IB-DDR (12x InfiniBand Double Data Rate)
- 12x IB-SDR (12x InfiniBand Single Data Rate)
- 1x IB-DDR (1x InfiniBand Double Data Rate)
- 1x IB-SDR (1x InfiniBand Single Data Rate).

The 12x IB-DDR (12x InfiniBand Double Data Rate) and 12x IB-SDR (12x InfiniBand Single Data Rate) optical links are used to connect a System z10 to either a System z10 or a System z9. They support a maximum link distance over IFB fiber optic cabling of 150 meters (492 feet) - three meters are reserved for intraserver connection.

The 12x IB-DDR and 12x IB-SDR optical links initialize at the single data rate (3.0 Gbps) and auto-negotiate to a higher speed (6.0 Gbps) if both ends of the link support the higher speed. For example, when a System z10 is connected to a System z9 using point-to-point IFB cabling, the link auto-negotiates to the highest common data rate – 3 Gbps. When a System z10 is connected to a System z10, the link auto-negotiates to the highest common data rate – 6 Gbps.

For a 12x IB-DDR optical link, the IFB cable is connected to an HCA2-O fanout card on the System z10. For a 12x IB-SDR optical link, the IFB cable is connected to an HCA1-O fanout card on the System z9. The IFB cable is a 50 micron OM3 multimode fiber optic cable with MPO connectors.

Both the HCA2-O fanout card and the HCA1-O fanout card contain two ports. Each port has an optical transmitter and receiver module.

The 1x IB-DDR and 1x IB-SDR optical links are used to connect a System z10 to another System z10. They support a maximum unrepeated distance is 10 kilometers (6.2 miles) and the maximum repeated distance is 100 kilometers (62 miles) when attached to a qualified Dense Wavelength Division Multiplexer (DWDM). The list of qualified DWDM vendors is available on Resource Link, (<http://www.ibm.com/servers/resourceLink>), located under the “Hardware products for server” on the **Library** page.)

The 1x IB-DDR optical link operates at a link data rate of 5 Gbps (Gigabits per second) when connected to a DWDM capable of DDR speed. The 1x IB-SDR optical link operates at a link data rate of 2.5 Gbps (Gigabits per second) when connected to a DWDM capable of SDR speed. The speed may be auto-negotiated if the attached DWDM is capable of operating at SDR or DDR speeds. This environment supports use of 9 micron single mode fiber optic cables with LC duplex connectors. The 1x IB-DDR and 1x IB-SDR cable is connected to an HCA2-O LR fanout card on the System z10 BC. HCA2-LR fanout card supports two ports for the 1x IB-DDR and 1x IB-SDR optical links.

Note: Note: The InfiniBand link data rates do not represent the performance of the link. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload. Specifically, with 12x InfiniBand coupling links, while the link data rate can be higher than that of ICB, the service times of coupling operations are greater, and the actual throughput is less.

When STP is enabled, IFB links can be used to transmit STP timekeeping information and can also be defined as timing-only links to other z10 BC servers, as well as z10 EC, z9 EC, z9 BC, z990, and z890 servers.

The CHPID type assigned to InfiniBand is CIB. Up to 16 CHPID type CIB can be defined to an HCA2-O fanout card, distributed across the two ports as needed. The ability to define up to 16 CHPIDs on an HCA2-O fanout allows physical coupling links to be shared by multiple sysplexes. For example, one CHPID can be directed to one CF, and another CHPID directed to another CF on the same target server, using the same port.

There is a maximum of 64 coupling CHPIDs for all link types (IFBs, ICs, ICB-4s, and active ISC-3s per server). z10 BC servers can support up to 12 IFB links. The maximum number of external coupling links combined (ICB-4, IFB, and ISC-3 links) cannot exceed 56 per server.

IC links

Internal coupling (IC) links are used for internal communication between coupling facilities defined in LPARs and z/OS images on the same server. IC link implementation is totally logical requiring no link hardware. However, a pair of CHPID numbers must be defined in the IOCDs for each IC connection. IC channels cannot be used for coupling connections to images in external systems.

IC links will have channel path CHPID type of ICP (Internal Coupling Peer). The rules that apply to the CHPID type ICP are the same as those that apply to CHPID type CFP (ISC-3 peer links), with the exception that the following functions are not supported:

- Service On/Off
- Reset I/O Interface
- Reset Error Thresholds
- Swap Channel Path
- Channel Diagnostic Monitor
- Repair/Verify (R/V)
- Configuration Manager Vital Product Data (VPD).

IC channels have improved coupling performance over ICB-4 and ISC-3 links. IC links also improve the reliability while reducing coupling cost. Up to 32 IC links can be defined on z10 BC; however, it is unusual to require more than one link (two CHPIDs type ICP).

Refer to “Internal coupling and HiperSockets channels” on page 41 for recommendations on CHPID usage.

Timing-only links

If you are configuring a timing network using STP, time synchronization may be required in configurations other than a Parallel Sysplex. For a server that is not part of a Parallel Sysplex, but required to be in the same Coordinated Timing Network (CTN), additional coupling links must be configured in order for the server to be configured in the CTN. These coupling links are called timing-only links. Use HCD to define Timing-only links and generate an STP control unit.

Coupling facility

The coupling facility provides shared storage and shared storage management functions for the Parallel Sysplex (for example, high speed caching, list processing, and locking functions). Applications running on z/OS images in the Parallel Sysplex define the shared structures used in the coupling facility.

PR/SM LPAR allows you to define the coupling facility, which is a special logical partition that runs Coupling Facility Control Code (CFCC). Coupling Facility Control Code is Licensed Internal Control Code (LICC). It is not an operating system. There are two versions of CFCC:

- ESA architecture (31-bit)
- z/Architecture (64-bit).

When the CFCC is loaded by using the LPAR coupling facility logical partition activation, the z/Architecture CFCC is always loaded. However, when CFCC is loaded into a coupling facility guest of z/VM, the ESA architecture or z/Architecture CFCC version is loaded based on how that guest is running.

At LPAR activation, CFCC automatically loads into the coupling facility LPAR from the Support Element hard disk. No initial program load (IPL) of an operating system is necessary or supported in the coupling facility LPAR.

CFCC runs in the coupling facility logical partition with minimal operator intervention. Operator activity is confined to the Operating System Messages task. PR/SM LPAR limits the hardware operator controls usually available for LPARs to avoid unnecessary operator activity. For more information, refer to *System z10 Processor Resource/Systems Manager Planning Guide*.

A coupling facility link provides the connectivity required for data sharing between the coupling facility and the CPCs directly attached to it. Coupling facility links are point-to-point connections that require a unique link definition at each end of the link.

CFCC considerations

To support migration from one coupling level to the next, you can run different levels of the coupling facility concurrently in different coupling facility LPARs on the same or different CPCs. Refer to “CFCC concurrent patch apply” for a description of how a CF patch or a new CFCC level can be applied to one coupling facility LPAR in a CPC while not affecting the code level of other coupling facility LPARs in the same CPC.

When migrating CF levels, the lock, list, and cache structure sizes may increase to support new functions. This adjustment can have an impact when the system allocates structures or copies structures from one coupling facility to another at different CFCC levels.

For any CFCC level upgrade, you should always run the CFSIZER tool which takes into account the amount of space needed for the current CFCC levels. The CFSIZER tool is available at <http://www.ibm.com/systems/support/z/cfsizer/>.

The content of each CFCC level is available at <http://www.ibm.com/systems/z/advantages/pso/cftable.html>.

CFCC concurrent patch apply

CFCC LICC maintenance and upgrades can be performed while the z/OS images connected to it continue to process work and without requiring a POR of the server on which the CF is located. This is done by installing the latest LICC changes onto the server, moving the structures on the CF to the second CF in the Parallel Sysplex, recycling the CF LPAR, and moving the structures back again. With CFCC concurrent patch apply, you can:

- Selectively apply the new patch to one of possibly several CFs running on a z10 BC. For example, if you have a CF that supports a test Parallel Sysplex and a CF that supports a production Parallel Sysplex on the same z10 BC, you now have the ability to apply a disruptive patch to only the test CF without affecting the production CF. After you have completed testing of the patch, it can be applied to the production CF as identified in the example.

- Allow all other LPARs on the z10 BC where a disruptive CFCC patch will be applied to continue to run without being impacted by the application of the disruptive CFCC patch.

This process does not change the characteristics of a concurrent CFCC patch, but does significantly enhance the availability characteristics of a disruptive CFCC patch by making it much less disruptive.

CFCC Level 16

CFCC Level 16 provides the following enhancements:

- Coupling Facility duplexing protocol enhancements provide faster service time when running System-managed CF structure duplexing by allowing one of the duplexing protocol exchanges to complete asynchronously. More benefits are seen as the distance between the CFs becomes larger, such as in a multisite Parallel Sysplex.
- CF subsidiary list notification enhancements provided to avoid false scheduling overhead for Shared Message Queue CF exploiters.
- Latch Fairness provides more consistent response time from CF requests.

CFCC Level 16 includes the support introduced in previous supported CFCC levels.

CFCC Level 14

CFCC Level 14 provides dispatcher and internal serialization mechanisms enhancements to improve the management of coupled workloads from all environments under certain circumstances.

CFCC Level 14 includes the support introduced with CFCC Level 13.

CFCC Level 13

CFCC level 13 provides Parallel Sysplex availability and performance enhancements. It provides changes that affect different software environments that run within a Parallel Sysplex. For example, DB2 data sharing is expected to see a performance improvement, especially for cast-out processing against very large DB2 group buffer pool structures.

CFCC Level 13 also includes features and functions from previous CFLEVELs.

CFCC Level 12

CFCC level 12 provides support for the following functions:

- **64-bit addressing**

The 64-bit addressing supports larger structure sizes and eliminates the 2 GB “control store” line in the coupling facility. With this support, the distinction between ‘control store’ and ‘non-control store’ (data storage) in the coupling facility is eliminated, and large central storage can be used for all coupling facility control and data objects.

- **48 internal tasks**

Up to 48 internal tasks for improved multiprocessing of coupling facility requests.

- **System-managed CF Structured duplexing (CF duplexing)**

CF duplexing is designed to provide a System z9 model, hardware assisted, easy-to-exploit mechanism for duplexing CF structure data. This provides a robust recovery mechanism for failures such as loss of single structure or CF, or loss of connectivity to a single CF, through rapid failover to the other structure instance of the duplex pair. Refer to “System-managed CF structure duplexing” on page 86 for more information.

Note: There is a substantial structure size growth (due to 64-bit support) when going from CFLEVEL 10 or 11 up to CFLEVEL 12. Structure re-sizing is required when migrating to CFLEVEL 12.

CFCC Level 11

CFCC Level 11 provides support for the following function:

- **System-managed CF structured duplexing (CF duplexing)**

CF duplexing is designed to provide an S/390 G5/G6 model, hardware assisted, easy-to-exploit mechanism for duplexing CF structure data. This provides a robust recovery mechanism for failures such as loss of single structure or CF, or loss of connectivity to a single CF, through rapid failover to the other structure instance of the duplex pair. Refer to “System-managed CF structure duplexing” on page 86 for more information.

Coupling connection considerations

There are several limits regarding coupling connections to be aware of when ordering and configuring these resources. Refer to Table 11 on page 76 for information on these link limits.

If individual link limits are exceeded, IOCP issues caution messages and HCD issues errors. Refer to *System z10 and z9 Stand-Alone Input/Output Configuration Program User's Guide* for details.

I/O configuration considerations

ICP IOCP supports ISC-3, ICB-4, IC, and IFB link definitions on z10 BC CPCs.

For z10 BC, HCD provides the following enhancements to support coupling facility definition:

- Controls for defining coupling facility channels. HCD also automatically generates the control unit and device definitions associated with CFP, CBP, CIB, and ICP channel paths, when they are connected to their respective peer or receiver channel paths. All IC channels paths must be connected.
- Controls for defining a logical partition as either a coupling facility or an operating system logical partition. HCD also allows the definition of the logical partition as both so its usage does not have to be pre-specified. This allows the flexibility of usage to be determined at logical partition activation. This way, if a partition is used one day as a coupling facility and the next day as a z/OS image logical partition, the I/O definitions do not need to change. Additionally, you must use these controls when defining a new logical partition in HCD.

IBM recommends that if you know a logical partition is used exclusively for ESA or exclusively for a coupling facility that you define it that way. This supplies the best HCD rule checking. They also recommend that you use the HCD when possible, to define the coupling facility channel configuration to the channel subsystem.

Server Time Protocol (STP)

Server Time Protocol (STP) provides the means for multiple System z10, System z9, z890, and z990 servers to maintain time synchronization with each other without using a Sysplex Timer. STP is designed to synchronize servers configured in a Parallel Sysplex or a sysplex without a coupling facility, as well as servers that are not in a sysplex.

Server Time Protocol is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of z10 BC server and CFs and presents a single view of time to Processor Resource/Systems Manager (PR/SM). STP uses a message-based protocol to transmit timekeeping information over externally defined coupling links between servers. The coupling links used to transport STP messages include ISC-3 links configured in peer mode, ICB-4 links, and IFB links. These links can be the same links already being used in a Parallel Sysplex for coupling facility communications.

By using the same links to exchange timekeeping information and coupling facility messages in a Parallel Sysplex, STP can scale with distance. Servers exchanging messages over short distance links, such as ICB-4 and IFB links, are designed to meet more stringent synchronization requirements than servers exchanging messages over long distance links, such as ISC-3 (distances up to 100 km), where the synchronization requirements are less stringent.

The STP design introduces a concept called Coordinated Timing Network (CTN). A Coordinated Timing Network (CTN) is a collection of servers and coupling facilities that are time synchronized to a time value called Coordinated Server Time. The concept of a Coordinated Timing Network fulfils two key goals:

- Concurrent migration from an existing External Time Reference (ETR) network to a time network using STP
- Capability of servers that cannot support STP to be synchronized in the same network as servers that support STP (z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890).

A CTN can be configured in two ways:

- **Mixed CTN** - Allows the coexistence of non-STP capable servers (z800 and z900 servers) in an External Time Reference (ETR) timing network with STP capable servers. In a Mixed CTN, the Sysplex Timer provides the timekeeping for the network. z800 and z900 servers are the only non-capable servers that can coexist in a Mixed CTN.
- **STP-only CTN** - Each server must be configured with same CTN ID. The HMC or SE provides the user interface for all time related functions, such as time initialization, time adjustment, and offset adjustment. The HMC or SE must also be used to initialize or modify the CTN ID and network configuration.

With STP you can:

- Initialize the time manually or by using an External Time Source (ETS). The ETS can be a dial-out time server or a connection to a Network Time Protocol (NTP) server so the Coordinated Server Time can be set to within 100 milliseconds of the ETS.
- Schedule periodic dial outs to a time service so that Coordinated Server Time may be gradually steered to an international time standard.

Note: If scheduled operations are set up to access the external time source and “Use NTP” or “Use NTP with pulse per second (PPS)” is selected on the ETS Configuration panel, the requests generated by scheduled operations will be ignored.

- Initialize Time Zone offset, Daylight Saving Time (DST) offset, and Leap seconds offset
- Schedule changes to Time Zone offset, Daylight Saving Time (DST) offset, and Leap Seconds offset. STP can automatically schedule Daylight Saving Time, based on the selected Time Zone

- Adjust Coordinated Server Time by up to +/- 60 seconds.
- Enhance the accuracy to an external time source by utilizing pulse per second (PPS) output from NTP servers. This allows you to set your time to an external time source to within 10 microseconds.
- Obtain external time source (ETS) by accessing the HMC that is running the NTP server. Using the HMC as an NTP server allows you to obtain your time source from the internet or another NTP server.
- Improve availability when power has failed for a single server (a server having the role of Preferred Time Server (PTS) and Current Time Server (CTS)) or when there is a site power outage in a multi-site configuration where the server having the role of PTS/CTS is installed.

If an Internal Battery Feature (IBF) is installed on your z10 BC server, STP can receive notification that power has failed and that the IBF is engaged. When STP receives this notification from a server that has the role of PTS/CTS, STP can automatically reassign the role of the Current Time Server (CTS) to the Backup Time Server (BTS), thus automating the recovery action and improving availability.

- Use an Application Programming Interface (API) to automate an STP CTN reconfiguration. If the Preferred Time Server (PTS) fails and the Backup Time Server (BTS) takes over as the Current Time Server (CTS), an API is available on the HMC so you can automate the reassignment of the PTS, BTS, and Arbiter roles. This improves availability by avoiding a single point of failure after the BTS has taken over as the Current Time Server (CTS).
- Save the STP configuration and time information across Power on Resets (POR) or power outages for a single or dual server STP-only CTN. This means you do not need to reinitialize the time or reassign the PTS/CTS role for a single server STP-only CTN or the PTS, BTS, or CTS roles for a dual server STP-only CTN across POR or power outage events.
- Generate z/OS messages and messages on the HMC when hardware events that affect the External Time Source (ETS) configured for an STP-only CTN occur.

The benefits of STP include:

- Allowing clock synchronization without requiring the Sysplex Timer and dedicated timer links. This reduces costs by eliminating Sysplex Timer maintenance costs, power costs, space requirements, and fiber optic infrastructure requirements.
- Supporting a multisite timing network of up to 100 km over fiber optic cabling, thus allowing a sysplex to span these distances. This overcomes the limitation of timer to timer links being supported only up to 40 km.
- Potentially reducing the cross-site connectivity required for a multisite Parallel Sysplex. Dedicated links are no longer required to transport timing information because STP and coupling facility messages may be transmitted over the same links.
- Supporting the configuration of different NTP servers for the Preferred Time Server (PTS) and the Backup Time Server (BTS), which improves the availability of NTP servers used as an external time source. Only the PTS or the BTS can be the Current Time Server in an STP-only CTN. If the PTS/CTS cannot access the NTP server or the pulse per second (PPS) signal from the NTP server, the BTS, if configured to a different NTP server, may be able to calculate the adjustment required and propagate it to the PTS/CTS. Then, the PTS/CTS will perform the necessary time adjustment steering.

System-managed CF structure duplexing

A set of architectural extensions to the Parallel Sysplex is provided for the support of system-managed coupling facility structure duplexing (CF duplexing) of coupling facility structures for high availability. All three structure types (cache structures, list structures, and locking structures) can be duplexed using this architecture.

Support for these extensions on z10 BC is concurrent with the entire System z family of servers in CFCC Level 12 and in z/OS. It also requires the appropriate level for the exploiter support of CF duplexing. Installing these releases and enabling this new function is designed to:

- Provide the necessary base for highly available coupling facility structure data through the redundancy of duplexing.
- Enhance Parallel Sysplex ease of use by reducing the complexity of CF structure recovery.
- Enable some installations to eliminate the requirement for standalone CFs in their Parallel Sysplex configuration.

For those CF structures that support use of CF duplexing, customers have the ability to dynamically enable (selectively by structure) or disable the use of CF duplexing.

The most visible change for CF duplexing is the requirement to connect coupling facilities to each other with coupling links. The required connectivity is bi-directional with a peer channel attached to each coupling facility for each remote CF connection. A single peer channel provides both the sender and receiver capabilities; therefore, only one physical link is required between each pair of coupling facilities. If redundancy is included for availability, then two peer mode links are required. However, this connectivity requirement does not necessarily imply any requirement for additional physical links. Peer mode channels can be shared between ICF partitions and local z/OS partitions, so existing links between servers can provide the connectivity between both:

- z/OS partitions and coupling facility images,
- Coupling facility images.

One of the benefits of CF duplexing is to hide coupling facility failures and structure failures and make total loss of coupling facility connectivity incidents transparent to the exploiters of the coupling facility. This is handled by:

- Shielding the active connectors to the structure from the observed failure condition so that they do not perform unnecessary recovery actions.
- Switching over to the structure instance that did not experience the failure.
- Reestablishing a new duplex copy of the structure at a specified time. This could be as quickly as when the coupling facility becomes available again, on a third coupling facility in the Parallel Sysplex, or when it is convenient for the customer.

System messages are generated as the structure falls back to simplex mode for monitoring and automation purposes. Until a new duplexed structure is established, the structure will operate in a simplex mode and may be recovered through whatever mechanism provided for structure recovery prior to the advent of CF duplexing.

As the two instances of a System-managed duplex structure get update requests, they must coordinate execution of the two commands to ensure that the updates are made consistently to both structures. Most read operations do not need to be duplexed.

z/OS operator commands display the status of the links for problem determination. In addition, the Resource Management Facility (RMF™) provides the performance management aspects about the CF-CF connectivity and the duplexed structures. Together, these enable the installation to manage and monitor the coupling facility configuration and new structure instances resulting from CF duplexing.

For more information on CF duplexing, you can refer to the technical white paper, *System-Managed CF Structure Duplexing* at the Parallel Sysplex website, <http://www.ibm.com/systems/z/psa/>.

Geographically Dispersed Parallel Sysplex (GDPS)

In e-business, two important objectives for survival are systems designed to provide continuous availability and near transparent disaster recovery (DR). Systems that are designed to deliver continuous availability combine the characteristics of high availability and near continuous operations to deliver high levels of service - targeted at 24 x 7.

To attain high levels of continuous availability and near-transparent disaster recovery, the solution should be based on geographical clusters and data mirroring. These technologies are the backbone of the GDPS solution. GDPS offers three different solutions based on the underlying mirroring technology:

- The GDPS solution, based on Peer-to-Peer Remote Copy (PPRC, recently renamed to IBM System Storage® Metro Mirror), is referred to as GDPS/PPRC.
- The GDPS solution based on Extended Remote Copy (XRC, recently renamed to IBM System Storage z/OS Global Mirror), is referred to as GDPS/XRC.
- The GDPS solution based on IBM System Storage Global Mirror is referred to as GDPS/Global Mirror.

GDPS is an integrated, automated application and data availability solution designed to provide the capability to manage the remote copy configuration and storage subsystem(s), automate Parallel Sysplex operational tasks, and perform failure recovery from a single point of control, thereby helping to improve application availability. GDPS is independent of the transaction manager (e.g., CICS® TS, IMS™, WebSphere®) or database manager (e.g., DB2, IMS, and VSAM) being used, and is enabled by means of key IBM technologies and architectures.

GDPS/PPRC

GDPS/PPRC is designed to manage and protect IT services by handling planned and unplanned exception conditions, and maintain data integrity across multiple volumes and storage subsystems. By managing both planned and unplanned exception conditions, GDPS/PPRC can help to maximize application availability and provide business continuity.

GDPS/PPRC is capable of the following attributes:

- Near continuous Availability solution
- Near transparent D/R solution
- Recovery Time Objective (RTO) less than an hour
- Recovery Point Objective (RPO) of zero (optional)

- Protects against localized area disasters (distance between sites limited to 100 km fiber)

Server Time Protocol (STP) allows you to have a GDPS/PPRC configuration across two sites up to 100 km apart. STP overcomes the limitations of the Sysplex Timer to Sysplex Timer links being supported only up to 40 km. STP is designed to eliminate the need for a third site to locate the second Sysplex Timer for installations where the multisite sysplex spans a distance greater than 40 km but less than 100 km.

GDPS/XRC

Extended Remote Copy (XRC, recently renamed to IBM TotalStorage z/OS Global Mirror) is a combined hardware and z/OS software asynchronous remote copy solution. GDPS/XRC includes automation to manage remote copy pairs and automates the process of recovering the production environment with limited manual intervention, including invocation of CBU. This provides the ability to perform a controlled site switch for a planned or unplanned site outage and gives significant value in reducing the duration of the recovery window and requiring less operator interaction.

GDPS/XRC is capable of the following attributes:

- Disaster recovery solution
- RTO between an hour to two hours
- RPO less than one minute
- Protects against localized as well as regional disasters (distance between sites is unlimited)
- Minimal remote copy performance impact.

GDPS/Global Mirror

Global Mirror enables a two-site disaster recovery and backup solution for z/OS and open systems environments. Using asynchronous technology, Global Mirror is designed to maintain a consistent and restartable copy of data at a remote site that can be located at virtually unlimited distances from the local site.

GDPS/Global Mirror is capable of the following attributes:

- Disaster recovery solution
- RTO between an hour to two hours
- RPO less than one minute
- Protects against regional disasters (distance between sites is unlimited)
- Minimal remote copy performance impact
- Support for z/OS and open data.

Intelligent Resource Director (IRD)

Intelligent Resource Director (IRD) is a function that optimizes your workload's resource utilization of the z10 BC across multiple logical partitions.

This extension strengthens key z10 BC and S/390 platform technologies, including z/OS Workload Manager, Processor Resource/Systems Manager (PR/SM) (logical partitioning hardware technology) and Parallel Sysplex Clustering technology. This powerful combination provides the ability to dynamically manage workloads within multiple logical operating system images executing on a single z10 BC server, as a

single large-scale computer resource, with dynamic workload management and physical resource balancing built into the native operating system and underlying hardware.

With IRD, z/OS WLM will exploit Parallel Sysplex technologies to monitor performance of workloads on multiple images against those workload goals. z/OS WLM will then interact with the PR/SM hipervisor, directing PR/SM to dynamically adjust the physical CPU and I/O resource allocation of the hardware across the multiple operating system instances, without requiring Parallel Sysplex data-sharing to achieve these benefits, and totally transparent to customer workload applications.

IRD not only combines PR/SM, z/OS WLM, and Parallel Sysplex for LPAR CPU management, but it also includes two additional z10 BC exclusives: Dynamic Channel Path Management (DCM), and I/O (Channel) Subsystem Priority to increase business productivity.

Through IRD technology extensions, the Parallel Sysplex will be able to dynamically change system image weights, reconfigure channels on the fly, and vary logical processors on and offline dynamically to maximize overall throughput across all of the system images to enable the most critical business application of highest priority to get the resources (CPU and I/O) it needs.

LPAR CPU management (clustering)

An LPAR cluster is the subset of the systems in a Parallel Sysplex that are running as logical partitions on the same server.

LPAR CPU management allows dynamic adjustment of processor resources across partitions in the same LPAR cluster. Through the z/OS WLM policy, installations specify the business importance and goals for their workloads. WLM will then manage these sets of logical partitions to provide the processor resources needed for the work to meet its goals based on business importance.

LPAR CPU management requires z/OS WLM goal mode and a coupling facility structure which contains critical status information enabling cross-partition management of CP and I/O resources.

LPAR CPU management can manage Linux on System z on an LPAR running on regular CPs, but not on IFLs.

I/O priority queuing (IOPQ)

I/O subsystem priority queuing extends the classic strengths of I/O priority queuing by addressing other challenges that are not currently handled by existing I/O priority schemes.

For example, prior to I/O subsystem priority queuing, discretionary work in one partition could dominate channels shared with business critical work in another partition. With this new function, z/OS WLM and the Hardware Management Console set priorities that will be used to give the business-critical work higher priority access to the channels. This in turn may allow customers that do not exploit MIF, in order to prevent such problems, to be able to do so now and may lead to reduced overall channel requirements. These new capabilities will help provide optimal workload management.

The range of I/O weights for each logical partition is set within the Hardware Management Console. WLM adjusts the I/O weights within this range. It can be a fixed range, in which WLM would play no part.

Dynamic channel path management (DCM)

This portion of IRD is a combination of hardware strengths and software flexibility. Paths can now be managed between the processor and the control units in the system. Dynamic Channel Path Management (DCM) enables the system to respond to ever changing channel requirements by moving channels from lesser used control units to more heavily used control units as needed. DCM can manage control units connected to ESCON channels.

When used with z/OS Workload Manager (z/OS WLM) in Goal Mode, z/OS WLM is able to direct Dynamic Channel Path Management to move channels to help business critical work achieve its goals. This also helps reduce the requirement for greater than 256 channels.

I/O priority queuing and Dynamic Channel Path Management (DCM) benefit the Parallel Sysplex environment, with increased benefit in a multi-image environment (Parallel Sysplex). Although Parallel Sysplex data sharing is not required for IRD, the benefits of combining the two are unsurpassed.

I/O priority queueing has no value in a single-system environment.

Table 12. IOPQ in a single-system environment

IRD function	Require CF?	Require goal mode?	Value in single-system cluster
LPAR CPU Mgmt	Yes	Yes	Little (Vary Logical CP)
DCM	Yes	No	Yes
IOPQ	No	No	No
Note: Both DCM and IOPQ do have more value with goal mode.			

Workload manager (WLM)

With the z10 BC, workload manager (WLM) provides industry leading partitioning and workload management. Maximum utilization of all system resources is enabled through dynamic, automatic allocation of processor, memory, and I/O resources across partitions based on real time workload demand and customer policy.

Workload manager on the z10 BC provides end-to-end management of transactions, from the web-browser to data storage then back to the web-browser. Workload manager can exploit Cisco routers and facilitate dynamic and automatic self-management of data based on business priorities.

Using IBM's discrete server technology with the z10 BC and z/OS, installations may take advantage of workload based pricing to further reduce the cost of computing as applications continue to grow by using:

- Software pricing based on what you define, not what capacity has been installed.
- Common pricing for many cross-platform products.
- License manager, which simplifies and centralizes via a standard licensing certificate to control software usage billing.

Workload based pricing is adopted by many tools vendors, and provides for 'rigid' management within a flexible system.

EAL5 certification

The z10 BC server received the Common Criteria Evaluation Assurance Level 5 (EAL5) certification level (Evaluation Assurance Level 5) for the security of its LPARs that run under the control of the Processor Resource/Systems Manager (PR/SM).

Chapter 7. Cryptography

System z10 EC offers a number of standard and optional hardware-based encryption features. These features include:

- CP Assist for Cryptographic Function (CPACF)
- Configurable Crypto Express2 and Crypto Express2-1P
- Configurable Crypto Express3 and Crypto Express3-1P.

CPACF delivers cryptographic support for Data Encryption Standard (DES), Triple Data Encryption Standard (TDES), Advanced Encryption Standard (AES), Secure Hash Algorithm (SHA), and Pseudo Random Number Generation (PRNG).

The Crypto Express2 feature (FC 0863), Crypto Express2-1P feature (FC 0870), Crypto Express3 feature (FC 0864), and Crypto Express3-1P feature (FC 0871) combine the functions of a coprocessor (for secure key encrypted transactions) and accelerator (for Secure Sockets Layer (SSL) modes into a single feature.

Support for CPACF is also available through the Integrated Cryptographic Service Facility (ICSF). ICSF is a component of z/OS that is designed to transparently use the CPACF and Crypto Express2, Crypto Express2-1P, Crypto Express3, or Crypto Express3-1P functions to balance the workload and satisfy the bandwidth requirements of the applications.

Products that include any of the cryptographic feature codes contain cryptographic functions that are subject to special export licensing requirements by the U.S. Department of Commerce. It is the your responsibility to understand and adhere to these regulations whenever moving, selling, or transferring these products.

The cryptographic features are eligible for export under License Exception ENC as retail items to all end users in all countries except the embargoed, subject to the usual customer screening. The dormant cards themselves, without the enabling software, are also eligible for export an NLR (No License Required) to all customers in all countries except the embargoed, subject to the usual screening.

CP Assist for Cryptographic Function (CPACF)

CPACF is available on the z10 EC server. The CPACF provides a set of symmetric cryptographic functions that focus on the encryption/decryption function of clear key operations for SSL, Virtual Private Network (VPN), and data storing applications not requiring FIPS 140-2 level 4 security. Each CPACF is shared between two processor units (PUs), which can be designated as various specialty engine types (CPs, IFLs, zIIPs, zAAPs).

The CPACF function is activated using a no-charge enablement feature (FC 3863) and offers the following support on every CFACF:

- Data Encryption Standard (DES)
- Triple data Encryption Standard (TDES)
- Advanced Encryption Standard (AES) for 128-bit, 192-bit, and 256-bit keys
- Secure Hash Algorithms: SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512
- Pseudo Random Number Generation (PRNG).

The DES, TDES, and AES functions use clear key values.

The DES, TDES, AES, and PRNG functions require enablement of the CPACF function (no charge FC 3863) for export control. The CPACF for SHA-1, SHA-224, SHA-256, SHA-384, and SHA-512 are shipped enabled.

MSA instructions for invoking CPACF function for DES, TDES, AES, PRNG, SHA-1, SHA-256, and SHA-512 are found in the *z/Architecture Principles of Operation*.

Protected key CPACF

When using CPACF for high performance data encryption, CPACF also helps to ensure that key material is not visible to applications or the operating systems. This support requires Crypto Express3.

The protected key function is supported by z/VM 5.4 or later.

Enablement and disablement of DEA key and AES key functions

Using the Customize Activation Profile task on the Support Element, you can enable the encrypt DEA key and encrypt AES key functions of the CPACF to import a clear key, then disable the encrypt DEA key and encrypt AES key functions to protect the CPACF from further imports. The CPACF feature must be installed to use the DEA key and AES key functions on the Support Element.

Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P

The Crypto Express2 feature (FC 0863) and Crypto Express3 feature (FC 0864) are designed to satisfy high-end server security requirements. Each Crypto Express2 or Crypto Express3 feature contains two adapters. The Crypto Express2-1P feature (FC 0870) and Crypto Express3-1P feature (FC 0871) are designed to satisfy small and mid-range server security requirements. Each Crypto Express2-1P or Crypto Express3-1P feature contains one adapter. Each adapter (referred to as a crypto) can be configured as a coprocessor or an accelerator.

A Crypto Express2, Crypto Express2-1P, Crypto Express3, or Crypto Express3-1P coprocessor is used for secure key encrypted transactions. This is the default configuration. A coprocessor:

- Supports highly secure cryptographic functions, use of secure encrypted key values, and user defined extensions (UDX)
- Supports secure and clear-key RSA operations.

A Crypto Express2, Crypto Express2-1P, Crypto Express3, or Crypto Express3-1P accelerator is used for SSL acceleration. An accelerator:

- Supports clear key RSA acceleration.
- Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol.

Although each Crypto Express2, Crypto Express2-1P, Crypto Express3, or Crypto Express3-1P feature occupies an I/O slot and each feature is assigned PCHID values (two PCHIDs for Crypto Express2 and Crypto Express3; one PCHID for Crypto Express2-1P and Crypto Express3-1P), they do not use Channel Path Identifiers (CHPIDs). They use cryptographic numbers.

All LPARs can have access to the Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P features, if the image activation profile

configures the Crypto to the LPAR. Cryptos can be dynamically added, moved, or deleted to or from LPARs without affecting the operating state of the LPAR.

Crypto Express2 and Crypto Express2-1P

Both Crypto Express2 and Crypto Express2-1P provide a PCI-X interface to the host.

The Crypto Express2 feature provides function that was previously offered by PCICA and PCIXCC features. Functionally, Crypto Express2 and Crypto Express2-1P are the same, other than Crypto Express2 containing two adapters and Crypto Express2-1P containing one adapter. The function supported includes:

- Consolidation and simplification using a single cryptographic feature
- Public key cryptographic functions
- Hardware acceleration for Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols
- User Defined Extension (UDX)
- 13-19 Personal Account Numbers (PANs)
- Secure (encrypted) keys for AES-128, AES-192, and AES-256
- 4096-bit key RSA management capability and clear key RSA acceleration.

- For Crypto Express2:

- Maximum number of features per server: 8
 - Number of cryptos per feature: 2
 - Maximum number of cryptos per server: 16
 - Number of domains per crypto: 16
 - Number of active LPARs per server: 60

- For Crypto Express2-1P:

- Maximum number of features per server: 8
 - Number of cryptos per feature: 1
 - Maximum number of cryptos per server: 8
 - Number of domains per crypto: 16
 - Number of active LPARs per server: 16

Crypto Express3 and Crypto Express3-1P

Crypto Express3 and Crypto Express3-1P, like Crypto Express2 and Crypto Express2-1P, are tamper-sensing, tamper-responding, programmable cryptographic cards designed to satisfy a System z10 BC server's high-end and mid-range security requirements.

Both Crypto Express3 and Crypto Express3-1P provide a PCI EXPRESS (PCIe) interface to the host. Dual processors operate in parallel to support the Common Cryptographic Architecture (CCA) with high reliability.

Functionally, Crypto Express3 and Crypto Express3-1P are the same, other than Crypto Express3 containing two adapters and Crypto Express3-1P containing one adapter.

In addition to supporting all the cryptographic functions available on Crypto Express2, Crypto Express3 includes:

- Improved performance for symmetric and asymmetric operations
- Dynamic power management to maximize RSA performance while keeping within temperature limits of the tamper-responding package

- Lock step checking of dual CPUs for enhanced error detection and fault isolation of cryptographic operations
- Updated cryptographic algorithms used in firmware loading with the TKE
- Cryptographic key exchanges between IBM CCA and non-CCA servers
- Secure remote key loading of encryption keys to ATMs, point of sale terminals (POS), and PIN entry devices
- PIN generation, verification, and translation functions
- Secure cryptographic key generation, installation, and distribution using both public and secret key cryptographic methods.

User-defined extensions

User-Defined Extensions to the Common Cryptographic Architecture (CCA) are supported on Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P. For unique customer applications, Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P will support the loading of customized cryptographic functions on z10 BC. Support is available through ICSF and the Cryptographic Support for z/OS and z/OS.e web deliverable. Under a special contract with IBM, as a Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P customer, you will gain the flexibility to define and load custom cryptographic functions yourself. This service offering can be requested by referring to the IBM “Cryptocards” website, then selecting the **Custom Programming** option.

The following is required for UDX support:

- One or more Crypto Express2, Crypto Express2-1P, Crypto Express3, or Crypto Express3-1P features
- A Hardware Management Console
- A TKE workstation, if the UDX requires access control point
- z/VM 5.3 or later for guest exploitation
- z/OS V1R6 or later
- Cryptographic support for z/OS and z/OS.e V1.6 or later web deliverable, and PTFs. Cryptographic support for z/OS.e V1R6/V1R7 web deliverable (no longer downloadable).

If you use a User Defined Extension (UDX) of the Common Cryptographic Architecture (CAA), you should contact your local UDX provider for an application update before ordering a new z10 BC or before migrating or activating your UDX application. Your UDX application must be migrated to CCA level 3.41 or higher before activating it on the z10 BC using Crypto Express2. For Crypto Express3, your UDX must be migrated from Crypto Express2.

See <http://www.ibm.com/security/cryptocards/> for more details.

Trusted Key Entry (TKE)

The Trusted Key Entry (TKE) workstation (FC 0839 / FC 0840) and the TKE 6.0 Licensed Internal Code (FC 0858) are optional features that provide a basic key management system for ICSF. TKE includes one Cryptographic coprocessor, which can be logged on with a passphrase or a logon key pair, to be used by security administrators for key identification, exchange, separation, update, backup, and management. Additionally, optional Smart Card Readers can be attached as a secure way to store keys. Up to 10 TKE workstations can be ordered.

Note: TKE workstation (FC 0840) will be available 1Q2010).

TKE 6.0 enhancements include:

- Grouping of domains across one or more host cryptographic coprocessors. This allows you to run domain-scoped commands on every domain in the group using one command or to run host cryptographic adapter scoped commands on every adapter in the group using one command.
- Stronger cryptography encryption for TKE inbound/outbound authentication. This includes:
 - Ability to issue certificates with 2048-bit key strength
 - Encryption of sensitive data sent between the TKE and Crypto Express3 host cryptographic coprocessors using a 256-bit AES key
 - Signing of transmission requests with a 2048-bit signature key, if the host coprocessor is a Crypto Express3 coprocessor
 - Signing of replies sent by a Crypto Express3 coprocessor on the host with a 4096-bit key.

The Trusted Key Entry (TKE) workstation supports four users:

- Auto-logged user, which provides tasks to perform basic operations
- Admin user, which provides setup and configuration tasks
- Auditor user, which provides tasks related to configuring and viewing the audited security events
- Service user, which provides tasks for servicing the TKE workstation.

The orderable TKE features are:

- TKE 6.0 code (FC 0858) **and** TKE Workstation (FC 0839 / FC 0840)
- TKE Smart Card Readers (FC 0885)
- TKE Additional Smart Cards (FC 0884)

The TKE workstations require the TKE 6.0 or TKE 5.3 code **and** the TKE unit that contains Ethernet capability and PCI-X card.

Note: The following features can only be carried forward to System z10 EC. They can no longer be ordered:

- TKE 5.3 code (FC 0854)
- Smart Card Readers (FC 0887)
- TKE Additional Smart Cards (FC 0888)

Trusted Key Entry (TKE) with Smart Card Readers

Support for an optional Smart Card Readers attached to the TKE 6.0 or TKE 5.3 workstation allows the use of smart cards that contain an embedded microprocessor and associated memory for key storage. Access to and use of confidential data on the smart card is protected by a user-defined Personal Identification Number (PIN).

Wizard for migrating cryptographic configuration data

A wizard on TKE is available to help you migrate Cryptographic configuration data from one Cryptographic coprocessor to a different Cryptographic coprocessor. Using the migration wizard will reduce the number of steps it takes to migrate data, therefore minimizing user errors and decreasing the duration of the migration.

You can migrate configuration data from:

- Crypto Express2 to Crypto Express2
- Crypto Express2 to Crypto Express3
- Crypto Express3 to Crypto Express3
- Crypto Express3 to Crypto Express2

The target Cryptographic coprocessor must have the same or greater capabilities as the Cryptographic coprocessor from which the data is migrating.

To locate the migration wizard and to view the criteria for migrating from Crypto Express3 or Crypto Express2, click **What's New** on the TKE Welcome page.

RMF monitoring

The Cryptographic Hardware Activity report provides information about the activities in Crypto Express2 and Crypto Express3 features. The request rate (number of requests per second) is reported per adapter. In addition, the utilization (how much of the interval the feature is busy) and the average execution time of all operations is reported.

FIPS certification

The tamper-resistant hardware security module, which is contained within the Crypto Express2, Crypto Express2-1P, Crypto Express3, and Crypto Express3-1P is designed to meet the FIPS 140-2 Level 4 security requirements for hardware security modules.

Remote loading of ATM and POS keys

Remote key loading refers to the process of loading Data Encryption Standard (DES) keys to Automated Teller Machines (ATMs) and Point of Sale (POS) devices from a central administrative site. These enhancements provide two important new features:

- Ability to load initial keys to an ATM or a POS device from a remote location
- Enhanced capabilities for exchanging keys with non-CCA cryptographic systems.

Chapter 8. Cabling

z10 BC utilizes Small Form Factor (SFF) connectors for ESCON (MT-RJ), FICON (LC duplex), ISC-3 (LC duplex), Gigabit Ethernet (LC duplex), and 10 Gigabit Ethernet (LC duplex). The speed of the link is determined by the architecture and ranges from 17 MBps (ESCON); 1 Gbps (Gigabit Ethernet); 1, 2, or 4 Gbps (FICON); to 10 Gbps (10 Gigabit Ethernet). Up to three generations of features can coexist on z10 BC. Each feature has its own unique requirements; transceiver, connector, unrepeat distance, and link loss budget.

Fiber optic cables for the z10 BC are available via IBM Site and Facilities Services.

IBM Site and Facilities Services has a comprehensive set of scalable solutions to address IBM cabling requirements, from product-level to enterprise-level for small, medium, and large enterprises. IBM Site and Facilities Services is designed to deliver convenient, packaged services to help reduce the complexity of planning, ordering, and installing fiber optic cables. The appropriate fiber cabling is selected based upon the product requirements and the installed fiber plant.

The services include:

- **IBM Facilities Cabling Services** — fiber transport system
- **IBM IT Facilities Assessment, Design, and Construction Services** — optimized airflow assessment for cabling

These services take into consideration the requirements for all of the protocols and media types supported on the System z10, System z9, and zSeries (for example, ESCON, FICON, Coupling Links, OSA) whether the focus is the data center, the Storage Area Network (SAN), the Local Area Network (LAN), or the end-to-end enterprise.

Under **IBM Facilities Cabling Services**, there is the option to provide IBM Fiber Transport System (FTS) trunking commodities (fiber optic trunk cables, fiber harnesses, panel-mount boxes) for connecting to the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890. IBM can reduce the cable clutter and cable bulk under the floor. An analysis of the channel configuration and any existing fiber optic cabling is performed to determine the required FTS trunking commodities. IBM can also help organize the entire enterprise. This option includes enterprise planning, new cables, fiber optic trunking commodities, installation, and documentation.

Under **IBM IT Facilities Assessment, Design, and Construction Services**, the Optimized Airflow Assessment for Cabling option provides you with a comprehensive review of your existing data center cabling infrastructure. This service provides an expert analysis of the overall cabling design required to help improve data center airflow for optimized cooling, and to facilitate operational efficiency through simplified change management.

For additional information on cabling, you can refer to any of the following:

- *System z10 Business Class Installation Manual for Physical Planning*
- Resource Link (<http://www.ibm.com/servers/resourcelink>), under **Services** from the navigation bar.

Fiber Quick Connect (FQC) for ESCON and FICON LX cabling

Fiber Quick Connect (FQC), an optional feature on z10 BC, is available for all ESCON (62.5 micron multimode fiber) and FICON LX (single-mode fiber) channels. FQC is designed to significantly reduce the amount of time required for on-site installation and setup of fiber optic cabling. FQC eases the addition of, moving of, and changes to ESCON and FICON LX fiber optic cables in the data center, and FQC may reduce fiber connection time by up to 80%.

FQC is for factory installation of IBM Facilities Cabling Services - Fiber Transport System (FTS) fiber harnesses for connection to channels in the I/O drawer. FTS fiber harnesses enable connection to FTS direct-attach fiber trunk cables from IBM Global Technology Services.

FQC, coupled with FTS, is a solution designed to help minimize disruptions and to isolate fiber cabling activities away from the active system as much as possible.

IBM provides the direct-attach trunk cables, patch panels, and Central Patching Location (CPL) hardware, as well as the planning and installation required to complete the total structured connectivity solution. For example, for ESCON, four trunks (each with 72 fiber pairs) can displace up to 240 fiber optic jumper cables, which is the maximum quantity of ESCON channels in one I/O drawer. This significantly reduces fiber optic jumper cable bulk.

On the CPL panels, you can select the connector to best meet your data center requirements. Small form factor connectors are available to help reduce the floor space required for patch panels.

Prior to the server arriving on-site, CPL planning and layout is done using the default CHannel Path IDentifier (CHPID) report and the documentation showing the CHPID layout and how the direct-attach harnesses are plugged.

Note: FQC supports all of the ESCON channels and all of the FICON LX channels in all of the I/O drawers of the server.

Cabling responsibilities

Fiber optic cables ordering, cable planning, labeling, and placement are the customer responsibilities for new installations and upgrades. Fiber optic conversion kits and Mode Conditioning Patch (MCP) cables are not orderable as features on a z10 BC. Representatives will not perform the fiber optic cabling tasks without a service contract.

The following tasks are required to be performed by the customer prior to machine installation:

- All fiber optic cable planning.
- All purchasing of correct, qualified, fiber cables.
- All installation of any required Mode Conditioning Patch (MCP) cables.
- All installation of any required Conversion Kits.
- All routing of fiber optic cables to correct floor cutouts for proper installation to server.
 - Use the Physical Channel Identifier (PCHID) report or the report from the Channel Path Identifier (CHPID) Mapping Tool to accurately route all cables.

- All labeling of fiber optic cables with PCHID numbers for proper installation to server.
 - Use the PCHID report or the report from the CHPID Mapping Tool to accurately label all cables.

Additional service charges may be incurred during the machine installation if the preceding cabling tasks are not accomplished as required.

Cable ordering

Fiber optic cables for the z10 BC are available from IBM Site and Facilities Services.

ICB-4 cables are available as features.

The following table lists the channel card feature codes and associated cabling information available on z10 BC. The features brought forward on an upgrade from z890 and z9 BC are also listed.

Table 13. Channel card feature codes and associated connector types and cable types

Feature Code	Feature Name	Connector Type	Cable Type
0163	HCA2-O	12x MPO	50 micron, OM3 12x IB-DDR
0168	HCA2-O LR	LC duplex	9 micron SM
0219	ISC-3	LC duplex	9 micron SM
2324 ¹	ESCON channel	MT-RJ	62.5 micron MM
2319 ¹	FICON Express LX	LC duplex	9 micron SM
2320	FICON Express SX	LC duplex	50, 62.5 micron MM
3318	FICON Express4-2C SX	LC duplex	9 micron SM
3319 ¹	FICON Express2 LX	LC duplex	9 micron SM
3320	FICON Express2 SX	LC duplex	50, 62.5 micron MM
3321 ¹	FICON Express4 10KM LX	LC duplex	9 micron SM
3322	FICON Express4 SX	LC duplex	50, 62.5 micron MM
3323 ¹	FICON Express4-2C 4KM LX	LC duplex	9 micron SM
3324 ¹	FICON Express4 4KM LX	LC duplex	9 micron SM
3325 ¹	FICON Express8 10KM LX	LC duplex	9 micron SM
3326	FICON Express8 SX	LC duplex	50, 62.5 micron MM
3362	OSA-Express3 GbE LX	LC duplex	9 micron SM ²
3363	OSA-Express3 GbE SX	LC duplex	50, 62.5 micron MM
3364	OSA-Express2 GbE LX	LC duplex	9 micron SM ²
3365	OSA-Express2 GbE SX	LC duplex	50, 62.5 micron MM
3366	OSA-Express2 1000BASE-T Ethernet	RJ-45	EIA/TIA Category 5 Unshielded Twisted Pair (UTP)
3367	OSA-Express3 1000BASE-T Ethernet	RJ-45	EIA/TIA Category 5 Unshielded Twisted Pair (UTP)
3368	OSA-Express2 10 GbE LR	SC duplex	9 micron SM
3369	OSA-Express3-2P 1000BASE-T Ethernet	RJ-45	EIA/TIA Category 5 Unshielded Twisted Pair (UTP)

Table 13. Channel card feature codes and associated connector types and cable types (continued)

Feature Code	Feature Name	Connector Type	Cable Type
3370	OSA-Express3 10 GbE LR	LC duplex	9 micron SM
3371	OSA-Express3 10 GbE SR	LC duplex	50, 62.5 micron MM
3373	OSA-Express3-2P GbE SX	LC duplex	50, 62.5 micron MM

Notes:

1. If this is an initial order and FQC is selected, the ESCON (FC 2324) and FICON (FC 2319, 3319, 3321, 3323, 3324, 3325) counts do not apply and are zeroed out.
2. Accommodates reuse of existing multimode fiber (50 or 62.5 micron) when used with a pair of mode conditioning patch (MCP) cables.

Refer to the **Services** section of Resource Link for additional information.

Cabling report

When the Fiber Quick Connect feature is ordered, a second part of the PCHID report is provided to document the connections between the ESCON and FICON LX channels and the MTP couplers. Figure 8 shows an example of the cabling portion of the report.

```

----- Fiber Trunking Section -----

Cage  Slot  F/C  Brkt      PCHID/Harn.Leg
A06B   4    3324  1F(1.3)L  1A0/2-3 1A1/2-4 1A2/2-5 1A3/2-6
A06B   5    3321  1F(1.3)L  1B0/3-1 1B1/3-2 1B2/3-3 1B3/3-4
A06B  10    2323  1R(6.10)R 1E0/6-1 1E1/6-2 1E2/6-3 1E3/6-4
                        1E4/6-5 1E5/6-6 1E6/7-1 1E7/7-2
                        1E8/7-3 1E9/7-4 1EA/7-5 1EB/7-6
                        1EC/8-1 1ED/8-2
A06B  11    2323  1R(6.10)R 1F0/8-4 1F1/8-5 1F2/8-6 1F3/9-1
                        1F4/9-2 1F5/9-3 1F6/9-4 1F7/9-5
                        1F8/9-6 1F9/A-1 1FA/A-2 1FB/A-3
                        1FC/A-4

Cage  Slot  F/C  Brkt      PCHID/Harn.Leg
A11B   4    3324  1F(4.6)L  220/5-3 221/5-4 222/5-5 223/5-6
A11B  11    2323  2R(6.10)R 270/8-4 271/8-5 272/8-6 273/9-1
                        274/9-2 275/9-3 276/9-4 277/9-5
                        278/9-6 279/A-1 27A/A-2 27B/A-3
                        27C/A-4

Legend:
A06B   I/O Drawer 1 in A frame
A11B   I/O Drawer 2 in A frame
3324   FICON Express4 4KM LX
2323   ESCON Channel 16 Ports
3321   FICON Express4 10KM LX

```

Figure 8. Cabling section of the PCHID report sample

The columns in this part of the report represent the following data:

Cage Displays the number of drawers installed and the location of the I/O and CPC drawers containing I/O. (CPC drawers are shown only if ICBs are part of the configuration.) In this example, there are two I/O drawers: one in drawer 1 (A06B) and one in drawer 2 (A11B).

Slot Displays the I/O slot where the harness is plugged.

F/C Displays the feature code of the channel card where the harness is plugged.

Brkt Displays the MTP bracket that the harness plugs into (**F** is the bracket in the front of the frame, **B** is the bracket in the back of the frame, **R** is the bracket to the right of the frame, **L** is the bracket to the left of the frame).

PCHID/Harn-Leg

Displays the PCHID number port harness is plugged into, the harness number based on the MTP coupler the harness is plugged to, the harness leg that is plugged into the port.

Chapter 9. Hardware Management Console and Support Element

The z10 BC include a Hardware Management Console and two internal Support Elements (SEs) located on the “A” frame. The second Support Element, the alternate SE, is standard on z10 BC and is configured the same as, and serves as an alternate to, the primary SE.

The Hardware Management Console operates as a closed platform which makes it easier to ensure the security and integrity. The purpose of the HMC is to provide a platform for the execution of the HMC application.

The HMC is configured with a firewall to limit network access in and out. By default, no external connections are allowed through the firewall. As objects are defined to the HMC application, the necessary firewall rules are added to allow for communications to and from these objects. Firewall rules are also added to allow external user connections, access by Product Engineering, and the customization of network settings to allow specific applications.

The Hardware Management Console communicates with each CPC through the CPC's Support Element. When tasks are performed at the Hardware Management Console, the commands are sent to one or more SEs which then issue commands to their CPCs. Commands can be sent to as many as all of the CPCs defined to the Hardware Management Console. One Hardware Management Console can control up to 100 CPCs and one CPC can be controlled by 32 Hardware Management Consoles. Refer to Figure 9 on page 106 for an example of a typical Hardware Management Console configuration.

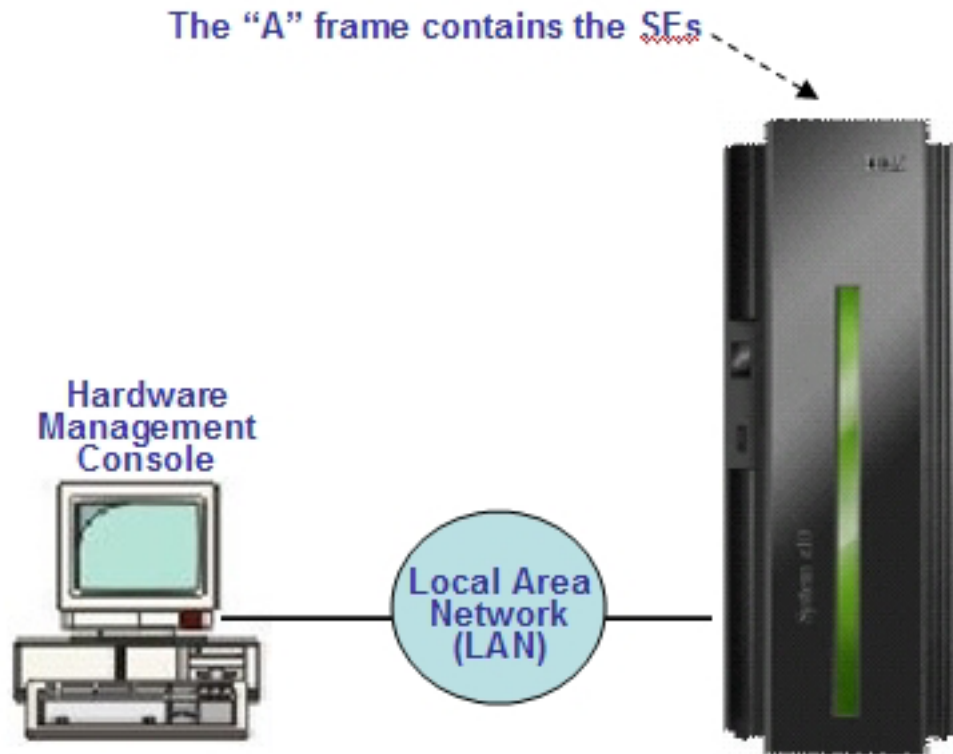


Figure 9. Hardware Management Console configuration

On z10 BC, CPCs configured to a Hardware Management Console are those CPCs whose internal SEs are:

- Attached by local area network (LAN) to the Hardware Management Console
- Defined to have the same domain name and domain password as the Hardware Management Console
- Defined in the Defined CPCs group (classic style view) or work pane (tree style view) at the Hardware Management Console.

The internal SEs for each CPC allows the Hardware Management Console to monitor the CPC by providing status information. Each internal SE provides the Hardware Management Console with operator controls for its associated CPC so that you can target operations:

- In parallel to multiple or all CPCs
- To a single CPC.

Hardware Management Console Application (HWMCA)

The Hardware Management Console Application (HWMCA) is a licensed software application installed on the Hardware Management Console. The application provides an easy-to-use object-oriented Graphical User Interface (GUI) you use to monitor and operate your CPCs. Starting the application makes the user interface available. You can directly manipulate objects displayed in the Hardware Management Console or Support Element workplace using a mouse or key combinations. The application begins whenever the console is powered-on or rebooted. For more detail about the Hardware Management Console and SE workplace, refer to the *System z Hardware Management Console Operations Guide*

and to the *System z10 Support Element Operations Guide*. You can also take an education course on Resource Link on *How to use the Hardware Management Console*.

Hardware Management Console and Support Element enhancements for z10 BC

This section highlights the significant changes and enhancements for the z10 BC Version Code 2.10.2. For more detailed information on these enhancements, see the *System z Hardware Management Console Operations Guide* and the *System z10 Support Element Operations Guide*.

Password prompt for disruptive actions

Use the **User Profiles** task, logged with user ID ACSADMIN, to determine whether you want to be prompted for a password for disruptive actions.

Fibre channel analysis

Use the **Fibre Channel Analyzer** task to identify fiber optic cabling issues in your Storage Area Network (SAN) fabric without contacting IBM service personnel. All FICON channel error information is forwarded to the HMC where it is analyzed to help detect and report the trends and thresholds for all FICON channels on System z10. This report shows an aggregate view of the data and can span multiple systems. This data includes information about the PCHID, CHPID, channel type, source link address, and destination link address of where the error occurred.

This applies to FICON channels exclusively (CHPID type FC).

To use this task, you must first enable the fibre channel analysis function located on the Customize Console Services task.

HiperSockets Network Traffic Analyzer Authorization task

Use the **Network Traffic Analyzer Authorization** task to select the level of authorization for HiperSockets network traffic analyzers for tracing one IQD channel or all IQD channels.

You can set the following:

- No traffic on any IQD channel for the selected server can be traced
- No traffic on the selected IQD channel can be traced
- All traffic on the selected IQD channel can be traced
- Customized traffic flow between selected logical partitions can be traced.

The **Network Traffic Analyzer Authorization** task is accessible by a user with the role of access administrator (default user ID, ACSADMIN).

Encrypt DEA key and Encrypt AES key functions

Use the **Customize Activation Profile** task on the Support Element to enable the encrypt DEA key and encrypt AES key functions of the CPACF to import a clear key and disable the encrypt DEA key and encrypt AES key functions to protect the CPACF from further imports.

Crypto Express3

Use the **Cryptographic Configuration** task, **Cryptographic Management** task, **Crypto Details** task, and the **View LPAR Cryptographic Controls** task to configure, manage, and view Crypto Express3 function.

New and revised tasks

Details about the HMC changes can be found in the *System z Hardware Management Console Operations Guide*. Details about the SE changes can be found in the *System z10 Support Element Operations Guide*.

Hardware Management Console and Support Element network connection

A local Hardware Management Console must be connected to its Support Elements through a Local Area Network (LAN). z10 BC use a dual ethernet (FC 0089) for the LAN wiring between the Hardware Management Console and the Support Elements. The necessary LAN adapters for the Support Elements and the Hardware Management Console may be specified as features on the system order.

Hardware Management Console (HMC) features and functions

Customization of the HMC or SE

You can use the Hardware Management Console workplace or Support Element workplace User Settings task to customize the presentation characteristics of the Hardware Management Console or Support Element. These customized settings can be saved to a diskette and used on other Hardware Management Consoles or Support Elements if desired. The User Settings task allows you to:

- Modify the default colors or use grey patterns instead of colors.
- Associate a color or pattern with any unacceptable status value you define to distinguish between types of exceptions.
- Change the background color of the Views area used to indicate exception and non-exception situations.
- Modify the default color associated with pending hardware or operating system messages.
- Enter the Hardware Management Console or Support Element TCP/IP address and domain name.

Status reporting

Each internal SE monitors the operation of its associated CPC and any CPC images running on the CPC and sends status to the Hardware Management Console for consolidation and exception processing.

Exception processing surfaces only those hardware status items you define as unacceptable to the running of your systems. You can define the conditions for exception processing at the Hardware Management Console or Support Element using the Details panel associated with each managed object.

- In the tree style user interface, the exceptions icon displays in the status bar if any managed object is in an unacceptable state. The exceptions icon also displays in the status column in the work pane next to the managed object that is in an unacceptable state.
- In the classic style user interface, the Hardware Management Console and Support Element display hardware status by using color (or grey patterns) to indicate acceptable or unacceptable status for objects. The default color change is from green (acceptable status) to red (unacceptable status). You can customize these colors (or patterns) using the Users Setting task.

Unacceptable status results in an exception situation that causes the color to change for the:

- Entire Views Area background.
- Object background in the Work Area for the object with the unacceptable status.
- Group object background in the Work Area for any group that the object with the unacceptable status is part of.

The default color change is from green (acceptable status) to red (unacceptable status). You can customize these colors (or patterns) using the Hardware Management Console workplace or Support Element workplace User Settings task.

Service Required state

The Service Required state indicates that the spare hardware shipped with the CPC has been depleted. When a part fails causing the use of the last redundant parts of that type, you now have just the required number of parts to keep the CPC running. This message is a reminder to you and the service representative that repairs should be made at the earliest possible time before addition.

The following conditions can cause a Service Required state:

- Loss of bulk power assembly (BPA)
- Primary SE loss of communications with the Alternate SE
- A PU is needed for sparing and no designated spare PU is available
- Insufficient spare PUs to support Capacity BackUp (CBU) or Disaster Recovery Assurance (if either feature is installed)
- Memory Sparing Threshold is reached
- High humidity is sensed inside the machine
- IML will cause CPC to go down or to an unacceptable condition
- Alternate SE is fenced due to automatic switchover.

Degrade indicator

The text “Degraded” indicates that, although the CPC is still operating, some hardware is not working. It displays on an object in the CPC group on the remote console, the Hardware Management Console, and on the SEs when:

- Loss of channels due to CPC hardware failure
- Loss of memory
- The drawer is no longer functioning
- Capacity BackUp (CBU) resources have expired
- Processor cycle time reduced due to temperature problem
- CPC was IMLed during cycle time reduction.

Hardware messages

The Hardware Management Console allows you to monitor the hardware messages from any CPC, CPC images, or any group of CPCs or CPC images configured to it. The Support Element allows you to monitor the hardware messages from its CPC or any CPC images configured to it.

Hardware messages present information about problems that are detected, suggest actions where possible, and aid in requesting service when appropriate. When a message has been handled, it is deleted from all the Hardware Management Console(s) and SE(s).

When hardware messages are pending for a hardware object or group of hardware objects:

- In the tree style user interface, the hardware message icon displays in the status bar if any managed object received a hardware message. The hardware message icon also displays in the status column in the work pane next to the specific managed object or objects that received a hardware message.
- In classic style user interface, the background of the object and its associated group turns blue (the default) and the Hardware Messages icon turns blue and flashes.

Operating system messages

Local operating systems and Coupling Facility Control Code (CFCC) running in a coupling facility partition can use the console integration facility of the hardware to send operator messages to be displayed by the Hardware Management Console or Support Element. The Hardware Management Console and Support Element allow you to monitor and respond to the operating system messages from any CPC image, coupling facility, or any group of CPC images configured to it.

For a coupling facility partition, Coupling Facility Control Code (CFCC) uses the console integration facility to display coupling facility messages and to accept Coupling Facility Control Code (CFCC) commands. The console integration facility, through the Operating System Messages task, provides the only interface for entering commands to an operating coupling facility partition.

When important operating system messages are pending for a hardware object or group of hardware objects:

- In the tree style user interface, the operating system message icon displays in the status bar if any managed object received an operating system message. The operating system message icon also displays in the status column in the work pane next to the specific managed object or objects that received an operating system message.
- In classic style user interface, the background of the object and its associated group turns cyan (the default) and the Operating System Messages icon turns cyan and flashes.

Problem analysis and reporting

Each primary SE monitors and analyzes problems detected on the associated CPC. For problems that are isolated to a single CPC, the results are reported to the Hardware Management Console as a hardware message. For those problems that potentially involve multiple CPCs, that problem data is sent to the Hardware Management Console, where data from multiple CPCs is analyzed and reported. The Hardware Management Console configured as a problem analysis focal point can perform:

- Problem analysis for FICON channel link errors of attached Support Elements.
- Problem analysis for ESCON, coupling facility, and Sysplex Timer link faults encountered by the CPCs configured to it.
- Service calls for all CPCs configured to it. Enabling the Hardware Management Console as a call home server identifies the Hardware Management Console as having a modem or LAN/Internet connection that all CPCs configured to it can use for placing service calls.

Virtual RETAIN

The Virtual RETAIN® function provides a way to capture problem data and place it in a temporary staging area on the Support Element hard disk for a problem that is to be called into IBM service. To ensure security and protection of the scan ring data, any hardware dump collected is encrypted before it is sent to RETAIN.

If RETAIN is not available, a hardware message is displayed for the Hardware Management Console, Support Element, and/or remote console user to instruct the customer to contact IBM Service to gather this staged problem data.

Licensed Internal Code (LIC)

Each Hardware Management Console and each SE has Licensed Internal Code (LIC) and is subject to periodic updates from IBM.

On systems with multiple Hardware Management Consoles, one of the Hardware Management Consoles should be configured as a LIC change management focal point. The Hardware Management Console configured can:

- Retrieve and distribute Licensed Internal Code updates for the Hardware Management Consoles remotely from IBM.
- Retrieve and distribute SE LIC updates to all the SEs of all the CPCs configured to the Hardware Management Console.

Remote I/O configuration and IOCDS management

Each CPC requires a definition of the I/O attached to it. The Hardware Configuration Definition (HCD) is a z/OS application that aids in the definition of all the I/O and aids in the distribution of the appropriate I/O definitions to the appropriate CPCs.

The Hardware Management Console configured as a change management focal point assists HCD in finding the names of all defined CPCs. A single HCD then distributes the appropriate IOCDS and IPL parameters to the various SEs of the CPCs defined to the same Hardware Management Console with change management capability.

Scheduled operations

The Hardware Management Console and Support Element provide support for scheduling the times and dates for automatic Licensed Internal Code (LIC) updates and backup of critical hard disk data for the Hardware Management Console, the CPCs configured to the Hardware Management Console, or the Support Element. You can accomplish this by using the Customize Scheduled Operations task.

For the Hardware Management Console, the Customize Scheduled Operations task, available from the HMC Management work pane (tree style view) or the Console Actions Work Area (classic style view), allows you to schedule the following LIC-related operations:

- Accept internal code changes
- Backup critical hard disk information
- Install concurrent code changes / Activate
- Remove internal code changes / Activate
- Retrieve internal code changes
- Retrieve internal code changes for defined CPCs
- Single step code changes retrieve and apply
- Transmit electronic service agent

- Transmit system availability data.

For the CPCs configured to the Hardware Management Console the Customize Scheduled Operations task, available from the Operational Customization tasks list, allows you to schedule the following LIC-related operations:

- Accept internal code changes
- Activate selected CPC
- Backup critical hard disk information
- Deactivate (Power off) selected CPC
- Install concurrent code changes / Activate
- Remove concurrent code changes / Activate
- Retrieve internal code changes
- Single step code changes retrieve and apply
- Transmit system availability data
- Access external time source
- Change LPAR weights.

For the Support Element the Customize Scheduled Operations task, available from the CPC Operational Customization tasks list, allows you to schedule the following LIC-related operations:

- Accept internal code changes
- Activate selected CPC
- Deactivate (Power off) selected CPC
- Install concurrent code changes / Activate
- Remove concurrent code changes / Activate
- Retrieve internal code changes
- Transmit system availability data
- Access external time source
- Change LPAR weights
- Activate or deactivate processor resources in an OOCOD record.

Remote Support Facility (RSF)

The Hardware Management Console provides Remote Support Facility (RSF) to aid in the service and maintenance of your system. RSF provides:

- Automatic or customer initiated call for service
- Automatic or customer downloading of the latest LIC change levels
- Automatic downloading of the latest phone list
- Support for records staged by Customer Initiated Upgrade (CIU)
- Support to enable Electronic Service Agent™ (Service Directory) to process operating system I/O error and software inventory data.

Remote Support Facility communicates with the IBM Service Support System using secure TCP/IP protocols. The communication may be through the enterprise LAN to the Internet (either directly or via a customer supplied SSL Proxy) or through the HMC-provided modem and customer supplied phone line. In both cases, the connection uses high grade SSL encryption, and is an outbound only connection.

Automation and API support

Application Programming Interfaces (APIs) on the Hardware Management Console and Support Element provide an end-user with the ability to view and manipulate managed objects.

The HMC supports both Common Information Model (CIM) and Simple Network Management Protocol (SNMP) as systems management APIs.

These APIs contain the ability to get/set a Hardware Management Console or Support Elements managed object's attributes, issue commands to be performed on a managed object from a local or remote application, and receive asynchronous event notifications. These APIs provide a mechanism to IBM, independent system management vendors, and an enterprise, to integrate with the Hardware Management Console Application (HWMCA).

For detailed information on the SNMP APIs, refer to *System z Application Programming Interfaces*. For detailed information on the CIM APIs, refer to *System z Common Information Model (CIM) Management Interface*.

CPC activation

Activating a CPC is an intelligent (LIC controlled) process that takes the CPC from its current state to a fully operational state. The activation may involve a power-on, power-on reset, and IPL, as necessary.

To activate a CPC, you must specify system activation information for each CPC configured to the Hardware Management Console.

You specify the CPC system activation information using activation profiles. Activation profiles allow you to define multiple power-on reset (POR) configurations for the same CPC. These profiles allow you to tailor the CPC resources (central processors, storage, and channels) to meet varying business and workload needs.

You use activation profiles to define PR/SM LPAR mode configurations. Activation profiles are stored on the SE hard disk so that they are available for future activate requests.

You can modify default activation profiles that are provided with the CPC and use them as templates to produce customized profiles to fit your processing needs.

There are four types of activation profiles you can produce:

- Reset - Used during power-on reset processing
- Load - Used to load an operating system
- Image - Used to define an logical partition.
- Group - Used to specify the capacity of a group of LPARs.

For PR/SM LPAR mode, you must define a reset profile and one Image profile for each logical partition.

NTP client/server support on the HMC

When the HMC has the NTP client installed and running, the HMC time can be continuously synchronized to an NTP server instead of synchronizing to a Support Element.

Also, when the HMC has the NTP client installed and running, the HMC can be configured to be used as an NTP server. This provides the ability for the Preferred Timer Server and Backup Time Server in an STP-only CTN to configure the external time source to use NTP with the HMC as the NTP server.

z/VM integrated systems management

z/VM integrated systems management from the HMC provides out-of-the-box integrated GUI-based basic management of z/VM guests. The HMC detects the

z/VM images. The z/VM integrated systems management function includes disk, network adaptor and memory management, guest activation and deactivation, and the display of guest status.

From the HMC, you can also:

- Dynamically determine if a directory manager is installed. If it is installed, the HMC allows any guests to be selected for management, whether it is running or not. It also allows for the defining, altering, and deleting of z/VM guests.

Enhanced z/VM systems management from the HMC allows selected virtual resources to be defined and managed, such as z/VM profiles, z/VM prototypes, z/VM virtual machines, and z/VM volume space.

- View and alter the Virtual Machine Resource Manager (VMRM) configuration and view the current VMRM measurement statistics.

Installation support for z/VM using the HMC

The z10 BC allows the installation of Linux on System z in a z/VM 5.4 or later virtual machine using the Hardware Management Console (HMC) DVD drive. This function does not require an external network connection between z/VM and the HMC. Instead, it utilizes the existing internal communication path between the Support Element and the HMC. Installing z/VM from the HMC DVD drive using the legacy support and the z/VM support, z/VM can be installed in an LPAR and both z/VM and Linux on System z can be installed in a virtual machine from the HMC DVD drive without requiring any external network setup or a connection between an LPAR and the HMC.

User authentication

You can configure the HMC to use a LDAP server to perform user ID and password authentication at logon time. The HMC still defines the user ID and the roles given to that user ID, but an enterprise directory (LDAP) server is used to authenticate the user. This eliminates the need to store the user ID's password locally.

This function allows the use of the current user ID/password policy for HMC user ID/passwords, and provides one centralized user ID/password control mechanism to help meet the user's corporate security guidelines.

Network protocols

The HMC for z10 BC uses a single network protocol, TCP/IP, when communicating with the Support Elements (SEs). This network protocol is used for both discovery and normal communications purposes.

The HMC supports IPv6 and IPv4 protocols within any customer network (for example, for remote access to the HMC user interface or for communication between HMCs and Support Elements). It can also perform electronic remote support requests to IBM service over an IPv6 or IPv4 network.

Customizable console date and time

The Customize Console Date and Time task uses more traditional time zone definitions rather than specific offsets from GMT. This allows for the automatic handling of special time zone characteristics such as daylight savings time.

System I/O configuration analyzer (SIOA)

The System I/O configuration analyzer allows the system hardware administrator access to the system's I/O configuration information from one place instead of

obtaining it from many separate applications. The analyzer makes it easier to manage I/O configurations, especially across multiple CPCs.

Network analysis tool for Support Element communications

A network analysis tool is available that allows you to verify that all required TCP/IP ports are supported from the HMC to the Support Element.

Instant messaging facility

An instant messaging facility is available that allows basic messaging capabilities between users of the HMC and the Support Element. It also allows messaging between local users and remote users using existing the HMC and Support Element interconnection protocols. The messaging capabilities include:

- Interactive chats between two partners using plain text
- Plain text broadcast message to all sessions on a selected console.

Screen capture function

The HMC allows you to capture full screens or specific windows of the HMC and save them as PNG, JPG, or GIF files. These files can then be viewed, copied to removable media, or deleted.

Call-home servers selection

You can select which HMCs can be designated as call-home servers for your Hardware Management Console.

User interface

The Hardware Management Console and Support Element allow you to choose the interface style in which you prefer to work:

- Tree style user interface
- Classic style user interface.

The tree style user interface is the default for Operator, Advanced Operator, Access Administrator, and System Programmer user roles. The classic user interface is the default for the Service Representative user role.

The **tree style** interface provides hierarchical views of system resources and tasks using drill-down and launch-in-context techniques to enable direct access to hardware resources and task management capabilities.

The **classic style** interface is the original user interface and has an object-oriented design. You can directly manipulate the objects (such as CPCs) that are defined and be aware of changes to hardware status as they are detected. There are several techniques for manipulating objects and tasks. One way to do this is to left-click an object to select it and double-click the task. An alternate method is the drag and drop technique, which involves using the mouse to pick up one or more objects, dragging them to a task, and then dropping them.

You can change from the tree style interface to the classic style using the User Settings task on the HMC.

Tree style user interface features

The following items are available when using the tree style user interface:

- **Tasks Index** node is available in the navigation pane. When selected, all tasks and their descriptions are listed in the work pane either in alphabetical order or by task category.
- **Expand all** and **collapse all** icons are available in the navigation pane and the task pad. The expand icon displays all the nodes in the navigation pane or all the tasks under each task group in the task pad. The collapse icon display only the main nodes in the navigation pane or the task groups in the task pad.
- **View** drop-down in the work pane table allows you to create a customized view in the work pane.

User authority

User authority for the HMC is determined by a user role that is associated with tasks. Each user role is actually a collection of authorizations. The HMC allows additional user roles to be defined to meet your needs. Refer to the *System z Hardware Management Console Operations Guide* for a list of predefined roles and details on how to define and customer user roles.

Available media

The HMC is no longer provided with a diskette drive. The available media is DVD-RAM, CD-ROM, and USB flash memory drive.

When using the Change Console Internal Code task to download internal code changes from the IBM Support System to removable media, you will be prompted for a user ID and password. It is the user ID and password used to access the IBM Support System. After the user ID and password are authenticated, the request will be permitted.

Security considerations

Because multiple Hardware Management Consoles and internal SEs require connection through a LAN, it is important to understand the use and capabilities enabled for each Hardware Management Console.

Hardware Management Consoles operate as peers with equal access to the CPCs configured to them. The SE for each CPC serializes command requests from Hardware Management Console Applications on a first come, first served basis. There is no guarantee of exclusive control across a sequence of commands sent from a single Hardware Management Console.

You should consider these security recommendations:

- Following installation of the CPC(s), Hardware Management Console(s), and SE(s) in your configuration, the access administrator should change the default logon passwords at the Hardware Management Console(s) and SE(s).
- Create a private LAN to interconnect the HMCs with the controlled SEs. Using a private LAN for your configuration offers several security, availability, and performance advantages as follows:
 - Direct access to the LAN is limited to the Hardware Management Console(s), SE(s), CPC(s), and control unit(s) attached to it. Outsiders cannot connect to it.
 - Traffic disruption due to temporary outages on the LAN is reduced, including disruptions caused by plugging in and powering-on new devices on the LAN (minor) to LAN adapters being run at the wrong speed (catastrophic).

- LAN traffic is minimized reducing the possibility of delays at the Hardware Management Console/SE user interface.
- Connect the HMCs to the enterprise LAN using the second LAN adapter in the HMC.
- Assign a unique domain name that includes all the CPCs controlled from one or more Hardware Management Consoles.
- Install one or more Hardware Management Consoles that have all of the CPCs you want to control defined to it.

Place at least one of these Hardware Management Consoles in the machine room near the CPCs that form its domain.

Use the following enable/disable controls to help you control access and provide focal point capabilities:

- Licensed Internal Code (LIC) update (change management focal point)
- Remote service support
- Remote customer access
- Remote service access
- Auto-answer of the modem.
- Physically secure the Hardware Management Console (keep it in a locked room).
- If a remote console is used for remote operations access, assign a secure logon password.
- Logoff each Hardware Management Console when it is not in use. The Hardware Management Console provides a status bar capable of displaying status colors (or grey patterns) to alert you when operator activity is needed, even when no one is logged on.
- Establish a limited list of objects and actions available to the operator.

Change management considerations

All Hardware Management Consoles are shipped with change management enabled. If you want to limit the number of Hardware Management Consoles that have change management capability such as LIC update control, I/O definition and remote IOCDS management capability using HCD, enable only those HMCs to be used as change management consoles. A least one Hardware Management Console in the domain must be enabled for change management.

Remote operations and remote access

Remote operations provides the ability to monitor or control a system, or group of systems, from a central or remote location. Remote capability creates a powerful tool for problem determination and diagnosis and operations assistance. Remote operations can save time and money while increasing the productivity of support staff. Technical expertise can be centralized, reducing the need for highly skilled personnel at remote locations.

Remote operations become increasingly important as:

- Data center operations and staff consolidate, with operations centers separate from those data centers
- Companies and their DP staffs merge
- World-wide operations become more common.

When considering remote operation of your z10 BC processor, there are two options available. You can choose one or both, based on your needs and configuration.

The first set of options deal with manual interaction and provide various methods of allowing a person to interact with the user interface. Manual control allows an operator to monitor and control the hardware components of the system using a hardware management console or a web browser.

A second set of options deal with machine interaction and provide methods of allowing a computer to interact with the consoles through an Application Program Interface (API). These automated interfaces allow a program to monitor and control the hardware components of the system. The automated interfaces are used by various automated products, including those from IBM and other vendors of other System Management products.

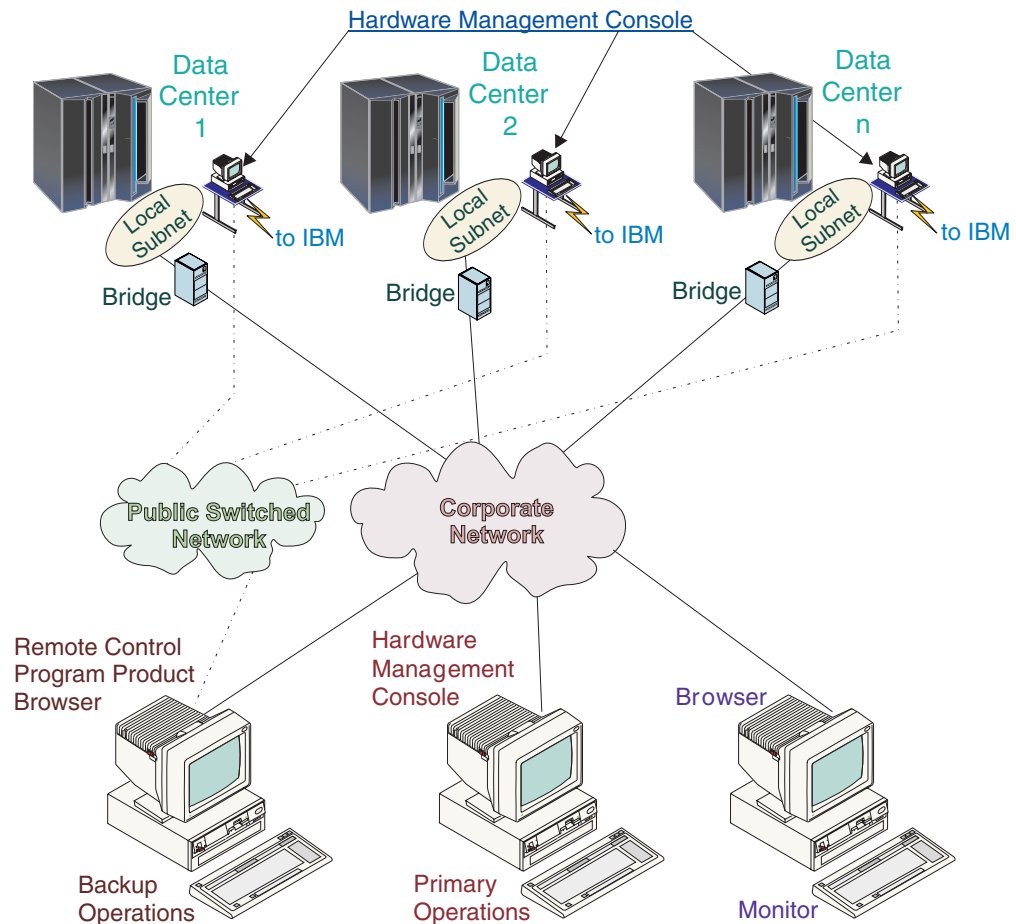


Figure 10. Remote operation example configuration

Remote manual operations

Remote manual operations use the same Graphical User Interface (GUI) used by a local HMC operator. There are two ways to perform remote manual operations, by:

- Using a remote HMC
- Using a web browser to connect to a local HMC.

The choice between a remote HMC and a web browser connected to a local HMC is determined by the scope of control needed. A remote HMC defines a specific set of managed objects that will be directly controlled by the remote HMC, while a web browser to a local HMC controls the same set of managed objects as the local HMC. An additional consideration is communications connectivity and speed. LAN

connectivity provides acceptable communications for either a remote HMC or web browser control of a local HMC but dialup connectivity is only acceptable for occasional web browser control.

Using a Hardware Management Console

A remote HMC gives the most complete set of functions because it is a complete Hardware Management Console; only the connection configuration is different from a local Hardware Management Console. As a complete HMC, it requires the same setup and maintenance as other HMCs. A remote HMC needs LAN TCP/IP connectivity to each Support Element to be managed. Therefore, any existing customer installed firewall between the remote HMC and its managed objects must permit communications between the HMC and SE. The remote HMC also requires connectivity to IBM or another HMC with connectivity to IBM for service and support.

Using a Web browser

Each HMC contains a Web server that can be configured to allow remote access for a specified set of users. When properly configured, an HMC can provide a remote user with access to all the functions of a local HMC except those that require physical access to the diskette or DVD media. The user interface on the remote HMC is the same as the local HMC and has the same constraints as the local HMC.

The Web browser can be connected to the local HMC using either a LAN TCP/IP connection or a switched, dial, network PPP TCP/IP connection. Both types connections can only use encrypted (HTTPS) protocols, as configured in the local HMC. If a PPP connection is used, the PPP password must be configured in the local HMC and in the remote browser system. Logon security for a web browser is provided by the local HMC user logon procedures. Certificates for secure communications are provided, and can be changed by the user.

Browser level is the responsibility of the customer and browser service or support and maintenance does not require connectivity to IBM.

Refer to the *System z Hardware Management Console Operations Guide*, SC24-6840 for Web browser requirements and information on getting ready to configure and use the Web Server and things to consider during your Web session.

Remote automated operations

As an alternative to manual operations, z10 BC allow a computer to interact with the consoles through a programmable interface, or API. The automated interface allows a program to monitor and control the hardware components of the system in the same way a human can monitor and control the system. The HMC APIs provide monitoring and control functions through TCP/IP SNMP to an HMC. These APIs provide the ability to get and set a managed object's attributes, issue commands, receive asynchronous notifications, and generate SNMP traps. For additional information about APIs, see the *System z Application Programming Interfaces*.

The automated interfaces are used by various automation products, including *Tivoli® System Automation for z/OS - Processor Operations*.

Chapter 10. Reliability, Availability, and Serviceability (RAS)

z10 BC reduces downtime by using standard features that provide high levels of Reliability, Availability, and Serviceability (RAS). This section lists the RAS functions available with the z10 BC:

Reliability

The features that provide a high level of reliability include the following:

- High-reliability technology components
- Parts integration to reduce the number of parts in the machine
- MCM System Run-In Test in manufacturing to reduce very early life fails
- To ensure data security, transmission of MCL files, restoration of backup files, and delivery of code loads via AROMs are digitally signed. In addition, any hardware dump collected from the Virtual RETAIN function is encrypted before being transmitted to RETAIN
- Common Criteria Evaluation Assurance Level 5 (EAL5) certification level for the security of its LPARs that run under the control of the Processor Resource/Systems Manager (PR/SM).

Availability

The functions that provide a high level of availability include the following:

Redundant I/O interconnect

Redundant I/O interconnect helps maintain critical connections to devices in the event of an MBA fanout card, HCA fanout card, STI cable, InfiniBand cable, or drawer failure by providing connectivity to the server I/O resources using a second path when available.

Plan ahead features

The Plan Ahead features allow you to order hardware that your current configuration will need in the future. Ordering ahead will avoid a disruptive hardware install in the future. The Plan Ahead features include: Plan Ahead Memory (FC 1991), Plan Ahead Memory Activation (FC 1992), Plan Ahead for Balanced Power (FC 3002) and Plan Ahead for I/O Expansion (FC 1999)

Plan Ahead Memory (FC 1991)

The Plan Ahead Memory feature (FC 1991) adds the necessary physical memory required to support target memory sizes. Therefore, it gives you the flexibility to activate memory to any logical size offered between the starting logical memory size and the target logical memory size. You can preplan future upgrades to be nondisruptive.

The Plan Ahead Memory feature is offered in 4 GB increments.

This feature is supported by z/OS and z/VM V5.4 or later.

Plan Ahead Memory Activation (FC 1992)

The Plan Ahead Memory Activation feature (FC 1992) allows you to activate a Plan Ahead Memory feature. One Plan Ahead Memory Activation is required for each Plan Ahead Memory feature.

Plan Ahead for Balanced Power (FC 3002)

The Plan Ahead for Balanced Power feature (FC 3002) allows you to order the maximum number of bulk power regulators (BPRs) on any configuration. This feature helps to ensure that your configuration will be in a balanced power environment if you intend to add I/O drawers to your server in the future.

Regardless of your configuration, all three BPR pairs will be shipped, installed, and activated.

Plan Ahead for I/O Expansion (FC 1999)

The Plan Ahead I/O Expansion feature (FC 1999) allows you to select the number of I/O drawers above the minimum number assigned by the configurator. This feature is intended for customers who plan to increase their I/O capability in the future and want to avoid the outage associated with an I/O drawer upgrade.

Enhanced driver maintenance

Licensed Internal Code (LIC) updates contribute to downtime during planned outages. The z10 BC can be configured to permit planned LIC updates to the server at specified driver sync points; that is, points in the maintenance process when LIC may be applied concurrently. A limited number of sync points exist throughout the life of the LIC level and once a sync point has passed, the LIC can no longer be applied concurrently. Availability is improved by taking advantage of these sync points.

Program directed re-IPL

Program directed re-IPL is designed to allow Linux on System z to re-IPL without operator intervention. Linux on System z can identify how to IPL was performed from the load device. Program directed re-IPL uses LIC to request a reload, using the same load parameters, from the same load device. Program directed re-IPL allows a Linux on System z running natively in an LPAR to execute a re-IPL.

z/OS V1R10 or later supports the program-directed IPL capability. The z/OS AutoIPL function allows the installation to specify IPL parameters to either IPL Stand-Alone Dump (SADMP), re-IPL z/OS, or both when the z/OS system requires a nonrestartable wait state to be loaded. z/OS also supports specification of IPL volumes and load parameters for IPLs that are to be performed.

Concurrent MBA fanout card and HCA fanout card hot-plug

The MBA fanout cards and HCA fanout cards reside on the front of a CPC drawer. In the event of an outage, an MBA fanout card or HCA fanout card, used for I/O, can be concurrently repaired using redundant I/O interconnect.

Processor unit (PU) sparing

In the rare case of a PU failure, the failed PUs characterization is dynamically and transparently reassigned to a spare PU. PUs not purchased by or reserved for a client will be available for transparent sparing and no repair action will be initiated as a result. Reserved, but currently unused, PUs will be available for sparing but

will generate a repair action. Where there are no available PUs, a failure will cause you to run degraded by one processor until a repair action can be completed.

Processor design

Each processor unit (PU) contains dual instruction and execution units that are used to verify internal operation and that operate simultaneously. Results are compared, and in the event of a miscompare, Instruction Retry is invoked. This design simplifies checking and virtually eliminates PU failures due to soft errors.

Support Element (SE)

The z10 BC has two Support Elements (SEs). In the event the primary SE fails, switchover to the alternate is usually handled automatically.

Hardware Management Console

One Hardware Management Console is required for system monitoring and operation of the CPC(s) configured to it. For high availability applications, it is recommended that you have at least two Hardware Management Consoles for your configuration to guarantee that the Hardware Management Console functions are available when needed.

The Hardware Management Console is concurrently maintainable with the operation of the CPCs configured to it. Having more than one Hardware Management Console provides continuous availability of Hardware Management Console functions, including the following:

- Hardware operator controls, hardware status monitoring, and hardware and operating system messages for all configured CPCs
- Capability to call for service
- Remote operations control
- Problem analysis.

Attaching to IBM service through the Internet

The z10 BC provides the ability to connect to IBM service using the Internet. The SSL connection is made from the HMC through the corporate network and firewall to IBM service using the Internet. This is an outbound connection only.

Hardware Management Console monitor system events

The Hardware Management Console monitor system events is available on z10 EC models. The Message and State Monitor facilities of the HMC can be enabled to send e-mail notification to a specified recipient whenever a specific message is received from either the hardware subsystem or an operating system, or when a CPC (hardware object) or a CPC image (Operating system object) changes from one "state" to another "state". The state of an object represents its current condition and functional capability, so a state change represents a change in functional capability that may require attention. Hardware and operating system messages are intended to keep the operator informed of conditions that may require attention. However, not all messages and not all state changes are important; only specific ones require attention and notification of a responsible person.

SAPs

z10 BC provides two base SAPs.

Application preservation

Application preservation is used in the case where a CP fails and there are no spares. The state of the failing CP is passed to another active CP where the operating system uses it to successfully resume the task in most cases without customer intervention.

Dynamic coupling facility dispatching

The dynamic coupling facility (CF) dispatching function helps enable continuous computing in the event of a coupling facility failure without requiring a standalone backup coupling facility. Enhanced dispatching algorithms enable you to define a backup coupling facility in a logical partition on your system. While this logical partition is in backup mode, although it is sharing resources with other logical partitions running other active workload, it uses very little processor resource. When the backup CF becomes active, only the resource necessary to provide coupling is allocated.

Error Correction Code (ECC)

Memory error checking and correction code detects and corrects single bit errors. Also, because of the memory structure design, errors due to a single memory chip failure are corrected.

Dynamic memory sparing

z10 BC does not contain spare memory DIMMs. Instead it has redundant memory distributed throughout its operational memory and these are used to bypass failing memory. Replacing memory cards may require power-off of the drawer, which is disruptive. The extensive use of redundant elements in the operational memory greatly minimizes the possibility of a failure that requires memory card replacement.

Memory scrubbing

Storage background scrubbing provides continuous monitoring of storage for the correction of detected faults before the storage is used.

Fixed HSA

Preplanning requirements are minimized by providing a fixed HSA (8 GB). A fixed HSA allows the maximum configuration capabilities to be exploited.

Dynamic additions to a channel subsystem and LPARs

You can dynamically add LPARs, LCSSs, subchannel sets, and logical CPs to an LPAR without preplanning.

You can dynamically update LPAR image profiles to support Crypto Express2 and Crypto Express3 without an outage to the LPAR. You can also dynamically delete or move Crypto Express2 and Crypto Express3 features from an LPAR.

LPAR dynamic storage reconfiguration

PR/SM LPAR storage reconfigurations can occur allowing nondisruptive add or removal to any partition with a cooperating guest. This capability removes the restriction of storage reconfigurations only being possible from an adjacent and above logical partition.

CICS subsystem storage protect

Subsystem storage protection and subspace group facility support, for use with CICS/ESA, prevents application software from overwriting CICS system software, control blocks, and address spaces.

Partial memory restart

In the event of a memory card failure, the system can be restarted with reduced memory capacity. Processing can be resumed until replacement memory is installed.

Dynamic I/O configuration

Dynamic I/O configuration enhances system availability without requiring a planned outage.

Dynamic I/O configuration allows you to add, delete, or modify the I/O definitions of channel paths, control units, and I/O devices in the CPC. You can also name previously reserved logical partitions and you can save the changes you made to the I/O configuration definitions and apply them to the active I/O Configuration Data Set (IOCDs).

Note: Dynamic I/O configuration is not available on a model with only ICFs or IFLs.

Dynamic I/O configuration requires z/OS or z/VM. Linux on System z, z/VSE, TPF, and z/TPF do not provide dynamic I/O configuration support.

When z/VM is controlling the I/O configuration, z/VM's dynamic I/O support is designed to handle all of the elements of the multiple Channel Subsystem facility for dynamic I/O configuration changes. To dynamically change the I/O configuration, one of two techniques can be employed:

- z/VM Control Program (CP) suite of interactive dynamic I/O commands
- HCM/HCD - configuration management tools.

Dynamic I/O configuration enhances system availability by supporting the dynamic addition, removal, or modification of channel paths, control units, I/O devices, and I/O configuration definitions to both hardware and software without requiring a planned outage.

Note: Dynamic I/O configuration is not available on a model with only ICFs or IFLs.

FICON cascaded directors

FICON cascaded directors allow a native FICON (FC) channel or a FICON Channel-to-Channel (CTC) to connect a server to a device or other server with two native FICON directors between them. This is only for a two-switch configuration.

FCP full-fabric connectivity

The FCP full-fabric connectivity supports multiple numbers of directors/switches that can be placed between the server and the FCP/SCSI device, thereby allowing many hops through a storage network for I/O connectivity.

Concurrent channel upgrade

It is possible to concurrently add ESCON, FICON, ISC-3, and OSA channels to an I/O drawer provided there are unused channel positions in the I/O drawer. In

addition, ICBs and IFBs and their associated cables can be added provided there are unused cable jack positions. This capability may help eliminate an outage to upgrade the channel configuration.

Dual power feeds

The power system offers dual primary (AC) power feeds. Each feed is electrically isolated and enables redundant power paths to each server. To take full advantage of the redundant power paths built into the z10 BC, it is necessary to provide dual electrical service to the server to minimize any outage due to a single-path power interruption.

Refer to *System z10 Business Class Installation Manual for Physical Planning* for more details about power feeds.

Redundant power and thermal subsystems

The DC power and thermal subsystems are designed with N +1 redundancy. Failure of a power or thermal component does not cause a system outage.

Oscillator (OSC) and External Time Reference (ETR) cards

There are two required OSC/ETR cards on the z10 BC. If an ETR function failure occurs on one card, it will automatically switch to the other card; the effectiveness of this action depends on the proper planning of the external ETR connections. Dynamic oscillator switchover is designed to detect the failure of the primary oscillator card and switch over to the backup, providing the clock signal to the server transparently.

Preferred Time Server and Backup Time Server

In an STP-only CTN configuration, it is required that at least one server is defined as the Preferred Time Server. It is also required that a Current Time Server is assigned. The Current Time Server (the Stratum 1 server) has the highest level of hierarchy in the STP-only CTN and has connectivity to the servers designated as Stratum 2 servers. If only a Preferred Time Server is defined in an STP-only CTN, it is assigned as the Current Time Server.

If there is a failure in the Preferred Time Server, synchronization fails if a backup configuration is not established. Therefore, it is highly recommended that a Backup Time Server is defined in an STP-only CTN. The Backup Time Server is normally a Stratum 2 server that has connectivity to the Preferred Time Server, as well as to all other Stratum 2 servers that are connected to the Preferred Time Server. By providing this connectivity, the Backup Server can easily take over as the Current Time Server if there is a failure with the Preferred Time Server or if a reconfiguration is planned. Therefore, the servers in the STP-only CTN can maintain synchronization.

Additionally, when the external time source for the STP-only CTN is configured to use NTP (with or without PPS), having the ETS configured on the Backup Time Server using different NTP server(s) provides continuous availability of NTP servers. In the event that the Preferred Time Server cannot access its configured NTP server(s), adjustments can be made using information from the Backup Time Server. This is achieved without reconfiguring the Backup Time Server as the Current Time Server.

Concurrent hardware maintenance

Concurrent maintenance enables the replacement of failed units concurrently with system operation. This enhances the processor availability by eliminating the need for system outage to effect the repair. Concurrent maintenance capability exists for the following elements:

- Power
- Thermal
- ESCON cards
- FICON Express cards
- OSA feature cards
- Crypto Express2 feature
- Crypto Express3 feature
- ISC-3 feature card
- ICB-4 STI cable
- MBA fanout card
- HCA2-O HCA2-O LR, HCA2-C fanout cards
- OSC/ETR card
- Hardware Management Console
- Support Element.

Concurrent Licensed Internal Code (LIC) patch

Concurrent LIC patch allows the activation of a patch concurrent with system operation thereby increasing the availability of the processor by reducing scheduled outage for LIC maintenance. This capability exists for code for the following elements:

- CP
- SAP
- Cage controllers
- LP
- CFCC
- Power
- Thermal
- ESCON channels
- FICON channels
- FCP channels
- OSA channels
- ISC-3 links
- ICB links
- IFB links
- IC links
- HiperSockets
- Hardware Management Console
- Support Element.

Notes:

1. OSA-Express channels always require CHPID vary off/vary on cycle to activate LIC patches.
2. Concurrent patch support is available for OSA-Express2 features.

Electronic Service Agent (Service Director)

Electronic Service Agent (Service Director™) will have I/O error data collected by a component integral to the operating system, forwarded from the operating system through a z10 BC Hardware Management Console, and then to an eService server

in IBM. Electronic Service Agent provides the analysis of the data and provides various users access to the data through a web browser interface.

Internal Battery Feature (IBF)

The Internal Battery Feature (IBF) provides backup input power. The feature is packaged internal to the machine. It can be used with a UPS to provide additional protection.

Redundant coupling links

Redundant coupling links (ISC-3s from different ISC-M cards, ICBs from different MBA fanout, and IFBs from different HCA fanout cards) can be configured between a processor and the coupling facility. This potentially removes a single point of failure for the processor's data sharing capability in the Parallel Sysplex environment.

Large page support

Large (1 megabyte (MB)) page support provides performance value primarily to long running memory access intensive applications.

Customer Initiated Upgrade (CIU)

Customer Initiated Upgrade (CIU) allows you to permanently increase processor or memory capacity. You can request these orders through the web using IBM Resource Link.

You can perform permanent upgrades while temporary capacity is active. This allows for quick conversion of temporary capacity to permanent capacity.

Capacity Upgrade on Demand (CUoD)

Capacity Upgrade on Demand provides the capability to permanently add CPs, ICFs, IFLs, zAAPs, zIIPs, SAPs, memory, and channels nondisruptively, eliminating the need for a scheduled outage. Installations who take advantage of the CUoD option may invoke the additional capacity nondisruptively.

On/Off Capacity on Demand (On/Off CoD)

When your business needs short term additional capacity, On/Off Capacity on Demand (On/Off CoD) is designed to deliver it. On/Off CoD is designed to temporarily turn on CPs, IFLs, ICFs, zAAPs, and SAPs.

Up to eight temporary records (CBU, CPE, and On/Off CoD) can be installed and activated at any given time.. You also have the flexibility of activating some of the resources on a given record. You do not have to activate the entire record. You also have the ability to add capacity and engines and extend the duration of the temporary upgrade concurrently, therefore eliminating the need for constant ordering of new temporary records for different customer scenarios.

You can order an On/Off CoD upgrade record using Resource Link.

Capacity Backup (CBU)

The z10 BC Capacity BackUp capability (temporary upgrade) enables enterprises to provide flexible, cost-effective Disaster Recovery on z10 BC.

Capacity for Planned Events (CPE)

Capacity for Planned Events (CPE) is designed to replace lost capacity within a customer's enterprise for planned down time events, such as system migration or relocation (for a data center move). This temporary upgrade is available for 3 days. You can order a CPE upgrade record using Resource Link.

Capacity provisioning

Capacity provisioning allows you to set up rules defining the circumstances under which additional capacity should be provisioned in order to fulfill a specific business need. The rules are based on criteria, such as: a specific application, the maximum additional capacity that should be activated, time and workload conditions.

This support provides a fast response to capacity changes and ensures sufficient processing power will be available with the least possible delay even if workloads fluctuate.

For more information, refer to the *z/OS MVS Capacity Provisioning Manager User's Guide*.

System-managed CF structure duplexing (CF duplexing)

A set of architectural extensions to the Parallel Sysplex is provided for the support of system managed CF structure duplexing (CF duplexing) of coupling facility structures for high availability.

Installing this software and microcode, and enabling this function is designed to:

- Provide the necessary base for highly available coupling facility structure data through the redundancy of duplexing.
- Enhance Parallel Sysplex ease of use by reducing the complexity of CF structure recover.
- Enable some installations to eliminate the requirement for standalone CFs in their Parallel Sysplex configuration.

Geographically Dispersed Parallel Sysplex (GDPS)

GDPS is a multisite solution that is designed to provide the capability to manage the remote copy configuration and storage subsystems, automate Parallel Sysplex operational tasks, and perform failure recovery from a single point of control, thereby helping to improve application availability.

GDPS/PPRC Hyperswap function is designed to broaden the continuous availability attributes of GDPS/PPRC by extending the parallel sysplex redundancy to disk subsystems.

Concurrent undo CBU

A prerequisite to executing this feature is that the customer or z/OS application must configure offline the processors that are being removed. So the best rule to follow is, "Configure offline the same logical processors that were configured online following the CBU activation." The concurrent undo CBU will require the following actions to configure the processors offline based on how it will be invoked:

- GDPS invoked Hardware Management Console/SE API
- Customer program invoked Hardware Management Console/SE API.

Notes:

1. As the user (or z/OS automation) deconfigures logical CPs, there is no guarantee that the logical CPs will remain in sequential numbering.
2. The SE panel will give no directions as to which CPs, ICFs, or IFLs to configure offline.

Fiber optic cabling

To serve the cabling needs of System z customers, IBM Site and Facilities Services has fiber optic cabling services available whether the requirements are product-level or enterprise-level. These services take into consideration the requirements for all of the protocols and media types supported on the System z10, System z9, and zSeries (for example, ESCON, FICON, coupling links, OSA) whether the focus is the data center, the Storage Area Network (SAN), the Local Area Network (LAN), or the end-to-end enterprise.

CHPID Mapping Tool

This tool provides a convenient interface to map hardware I/O ports on order to your CHPID definitions. An availability option automatically assigns PCHIDs to your CHPID definitions to minimize connecting critical paths to a single points of failure. This is recommended for all new z10 BC hardware builds or for upgrades from a z9 BC processor to a z10 BC, as well as for making changes to an already installed machine after hardware upgrades that change or increase the number of channels.

Multipath initial program load

z/OS on System z10 BC allows the system to attempt an IPL on alternate paths, if available, if I/O errors occur during the IPL. The system will attempt the IPL on an alternate path until all paths have been attempted or until the IPL is successful. This function increases the availability by allowing an IPL to complete using alternate paths and eliminates the need for manual problem determination when attempting an IPL.

This function is applicable for all FICON features with CHPID type FC and all ESCON features with CHPID type CNC.

Point-to-point SMP network

For System z10 BC, point-to-point SMP network provides growth paths up to a 10 engine system, where each of the 10 PUs have full access to all system resources, specially memory and I/O.

System-initiated CHPID reconfiguration

This function allows you to submit one request to all operating systems to configure offline or configure online all the CSS.CHPIDs associated with a particular chid. It reduces the duration of a repair action when an ESCON or FICON channel; an OSA port; or an ISC-3, ICB-4, or IFB link is shared across logical partitions (LPARs).

Link aggregation support

Link aggregation (trunking) is designed to allow you to combine multiple physical OSA ports of the same type into a single logical link. You can have up to eight OSA ports in one virtual switch. This increases bandwidth and permits nondisruptive

failover in the event that a port becomes unavailable. This function dedicates an OSA port to the z/VM 5.3 or later operating system for link aggregation under z/VM Virtual Switch-controlled link aggregation.

This support also provides dynamic add/remove of OSA ports and full-duplex mode (send and receive).

This support applies to OSA-Express2 and OSA-Express3.

Network Traffic Analyzer Trace facility

The Network Traffic Analyzer Trace facility is a diagnostic tool used to capture data as it enters or leaves an OSA adapter or Hipersockets channel for an attached host.

For OSA adapters, this facility is controlled and formatted by the z/OS Communications Server; however, the traced data is collected in the OSA at the network port.

For HiperSockets channels, the Support Element sets up authorization to allow tracing on selected HiperSockets. Traced data can be collected in a Linux partition; then tcpdump tools can be used to format and analyze the data.

For OSA adapters and HiperSockets Layer 2 devices, because the data is collected at the Ethernet frame level, you can trace the MAC headers for packets. For OSA adapters and HiperSockets channels, you can trace ARP packets, SNA packets, and packets being sent to and from other users sharing the OSA adapter or Hipersockets channel, including other TCP/IP stacks, Linux on System z users, and z/VM guest exploitation.

The Network Traffic Analyzer Trace facility supports OSA-Express2, OSA-Express3, and Hipersockets.

QDIO diagnostic synchronization

Queued Direct I/O (QDIO) diagnostic synchronization provides the ability to coordinate and simultaneously capture software (z/OS) and hardware (OSA) traces. This function allows the host operating system to signal the OSA feature to stop traces and allows the operator to capture both the hardware and software traces at the same time. You can specify an optional filter that alters what type of diagnostic data is collected by the OSA adapter. This filtering reduces the overall amount of diagnostic data collected and therefore decreases the likelihood that pertinent data is lost.

This support applies to OSA-Express2 and OSA-Express3.

FICON purge path extended

The FICON purge path error-recovery function is used in FICON problem determination. The FICON purge path error-recovery function can transfer error-related data and statistics between the channel and entry switch, and from the control unit and its entry switch to the host operating system.

FICON Express8 and FICON Express4 pluggable optics for individual servicing

The FICON Express8 and FICON Express4 features have small form factor (SFF) pluggable 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.

CICS subspace group facility

z10 BC CPCs provide support for the subspace group facility that can enhance the data integrity and reliability of application server subsystems, such as Customer Information Control System Transaction Server (CICS TS), reducing application failures, service outages, and incorrect modification of critical business data.

Dynamic channel path management

Dynamic channel path management (DCM) enables the system to respond to changing channel requirements by moving channels from lesser used control units to more heavily used control units as needed. DCM can manage control units connected to ESCON channels.

When used with z/OS Workload Manager (z/OS WLM) in goal mode, z/OS WLM can direct Dynamic Channel Path Management to move channels to help business critical work achieve its goals. This also helps reduce the requirement for greater than 256 channels.

Serviceability

The features that provide a high level of serviceability include the following:

- Automatic error detection and fault isolation concurrent with system operation.
- Automatic remote support capability.
- High degree of concurrent maintenance capability in hardware and code.
- Multiple Channel Swap - an enhancement for channel problem determination allowing up to 16 channels to be swapped.
- Status Panel showing status of N+1 power system.

Appendix A. System z10 BC Version 2.10.2 purpose and description

This appendix contains detailed information about Version 2.10.2 licensed internal code.

Preventative Service Planning (PSP) bucket considerations

Use IBM Service Link or contact your IBM representative to obtain a current copy of the 2098DEVICE bucket applicable to your environment. The PSP bucket contains corequisite software and hardware planning information that applies to various operating system environments. This includes Authorized Program Analysis Reports (APARS), Program Temporary Fixes (PTFs), and Licensed Internal Code (LIC) product patches.

Software corequisites

See the appropriate 2098DEVICE Preventative Service Planning (PSP) buckets subset ID for APAR and PTF information for the System z10 BC.

Table 14. Software corequisites

Software	PSP bucket subset ID
z/OS	2098ZOS
z/VM	2098ZVM
z/VSE	2098ZVSE

Engineering change (EC) considerations

Version Code 2.10.2 for z10 BC includes the following Central Processor Complex (CPC) and Hardware Management Console (HMC) Licensed Internal Code (LIC), engineering change (EC) and Microcode Load (MCL) levels:

- CPC level: EC N24409 + MCLs
- HMC Level: EC N24415 + MCLs

To verify that the enhancements described in this document apply to your system, display the LIC EC levels running on the CPC and the HMC.

CPC EC N24409 + MCLs

From the HMC using the tree view, you can display the LIC EC and MCL level of the system's CPC as follows:

1. From the navigation pane, select **Tasks Index**.
2. Scroll down the Tasks Index work pane and select **System Information**. The **Target Object Selection** window displays.
3. Select the object and click **OK**. The **System Information** window displays
4. Verify that the EC level is in this list.

HMC EC N24415 + MCLs

From the HMC using the tree view, you can display the LIC EC and MCL level of the system's HMC as follows:

1. From the navigation pane, select **Tasks Index**.

2. Scroll down the Tasks Index work pane and select **View Console Information**. The **View Console Information** window displays.
3. Verify that the EC level is in this list.

Miscellaneous lower level ECs included in Version 2.10.2

The following table provides miscellaneous changes included in CPC system code EC N24409 with Hardware Management Console system code EC N24415.

Table 15. ECs included in Version 2.10.2

EC number	Name
H25080	Backup-DVD New Build & MES Upgrade
N24422	SUL-DVD Driver 79
H25078	Security-Log DVD
G40983A	HMC/SE DIAGS CDR
J14654	TKE Backup-DVD New Build & MES Upgrade
J14860	SE Upgrade Data DVD MES Only
N10984	HMC DIAGS CDR
G43459	SE DIAGS CDR
J10269	Legacy SE Restore Diskettes for use with HMC Version 2.9.0 or higher
G40571	Blank Formatted Diskette
G40676 G40676A	STP Enablement Code (FC 1021)
G42645	HWMCA PC BIOS Update CDR Version 2.9.X
G35949	HWMCA Upgrade Data DVD MES Only
G42891	HWMCA Upgrade Data DVD

Appendix B. Resource Link

Resource Link is a customized web-based solution that provides everything you need to plan for, install, and maintain IBM System z10, IBM System z9, IBM eServer™ zSeries and IBM S/390 servers and associated software.

You can access Resource Link at <http://www.ibm.com/servers/resource link>.

Resource Link content areas include:

- **Personal Folders**
Organize site content according to your personal interests by creating folders for shortcuts to subscriptions. Subscribing to pages or individual documents means you will be notified of updates and changes.
- **Planning**
Interactive planning provides a streamlined plan for the installation of a system using online planning information tailored for your system.
- **Education**
A web-based multimedia education provides product courses that can be used for training or refreshing skills.
- **Library**
Product documentation that can be viewed, printed, or downloaded.
- **Fixes**
Interactive tools allow you to request, receive, and install system upgrades.
- **Problem Solving**
Resources to help you investigate and solve specific hardware and software problems.
- **Services**
Support for services such as installation, migration, networking, planning and relocating servers, Fiber Cabling, System z Application Programming Interfaces (APIs), and links to IBM software support.
- **Tools**
Information about tools such as machine information, CHPID mapping, coupling facility structure sizer, and links to software tools.
- **Customer Initiated Upgrade (CIU)**
A web-based application that allows you to download licensed internal code (LIC) to permanently upgrade processors and memory. You can also temporarily add processor capacity using the On/Off Capacity on Demand (On/Off CoD), Capacity for Planned Events (CPE), and Capacity Backup (CBU) features.

Resource Link Functions

Resource Link contains the following additional functions:

- **Customized planning aids** - Prepares you for the arrival and installation of your new System z10 BC hardware. To use Customized Planning Aids you need a valid order number and a Configuration Control Number (CCN), both available from your IBM Sales Representative.
- **CHPID Mapping Tool** - Downloadable from Resource Link, this tool allows you to map hardware I/O ports on order to your IOCP CHPID definitions. An availability option automatically maps hardware I/O ports to CHPIDs minimizing single points of failure.

- **Machine information** - Provides detailed information about your System z10, System z9, or zSeries machine including information about the last time your machine called home, service states, and hardware status.
- **Power Estimation Tool** - Allows you to estimate the power consumption of a specific System z10 or System z9 machine model and its associated configuration.
- **WWPN Prediction Tool** - Assists you in preplanning and setting up your Storage Area Networks (SANs) environment prior to the installation of your System z10 server. Therefore, you can be up and running much faster after the server is installed. This tool applies to all FICON channels defined as CHPID type FCP (for communication with SCSI devices).

Appendix C. Capacity upgrades

The z10 BC is designed to support concurrent upgrades that provide additional capacity with no server outage. The Capacity on Demand offerings provide permanent and temporary upgrades. All the upgrades are delivered by Licensed Internal Code Configuration Control (LICCC).

Licensed Internal Code Configuration Control (LICCC) provides for processor or memory upgrade with no hardware changes by enabling the activation of additional installed capacity. Concurrent upgrades using LICCC can be done for:

- CPs, SAPs, ICFs, IFLs, zIIPs, and zAAPs - requires available unused PUs in the installed drawer
- Memory - requires available capacity on installed memory cards
- Channel cards - requires available ports on channel cards.

You can order permanent upgrades using the Customer Initiated Upgrade (CIU) application through Resource Link or calling your IBM sales representative.

There are three type of temporary upgrades available on z10 BC. The offerings providing these upgrades are: On/Off Capacity on Demand (On/Off CoD), Capacity Backup (CBU), or Capacity for Planned Events (CPE). You can order a CPE and CBU temporary upgrade using the CIU application through Resource Link or calling your IBM sales representative. You can order an On/Off CoD temporary upgrade **only** using the CIU application through Resource Link.

Each Capacity on Demand offering is available through an IBM contract. You must order the Online CoD Buying feature (FC 9900) to enable using Resource Link to order capacity upgrades. Refer to the *System z10 Capacity on Demand User's Guide* for details.

Permanent upgrades

When using the CIU application through Resource Link to order a permanent upgrade, you can increase model capacity, add specialty engines (ICFs, IFLs, zAAPs, zIIPs, and SAPs), add memory, activate unassigned model capacity or IFLs, and deactivate activated model capacity or IFLs.

You can perform permanent upgrades while temporary capacity is active. This allows for quick conversion of temporary capacity to permanent capacity.

When calling your IBM sales representative to order a permanent upgrade (referred to as Capacity Upgrade on Demand (CUoD)), you can increase model capacity, add specialty engines (ICFs, IFLs, zAAPs, zIIPs, and SAPs), add memory, activate unassigned model capacity or IFLs, deactivate activated model capacity or IFLs, activate channels, activate crypto, and perform recharacterization.

Refer to the *System z10 Capacity on Demand User's Guide* for more information.

Temporary upgrades

Using On/Off CoD, CBU, or CPE, you can increase model capacity and add specialty engines (ICFs, IFLs, zAAPs, zIIPs, and SAPs).

Characteristics of temporary upgrades include:

- **Permanent upgrade while temporary capacity is active** - You can add permanent processor or memory capacity while temporary On/Off CoD, CBU, or CPE records are active. This allows for quick conversion of temporary capacity to permanent capacity.
- **Multiple records can be simultaneously active** - Up to eight records (On/Off CoD, CBU, and CPE) can be active at any given time. However, only one On/Off CoD record can be active at any given time.
- **Store LICCC records in an unactivated state** - Up to 200 records (On/Off CoD, CBU, and CPE) can be staged on the Support Element at any given time. This provides greater flexibility to quickly enable needed temporary capacity.
- **Automatic deactivation** - When a record expires, the resource is automatically deactivated. However, the record will not be deactivated if it means removing a dedicated engine or the last of that engine type.
- **Partial activation** - You do not have to activate the entire record. You can choose partial activation of resources up to the maximum you ordered.

On/Off Capacity on Demand (On/Off CoD)

On/Off Capacity on Demand (On/Off CoD) is designed to satisfy your need for short term additional capacity. On/Off CoD allows you to temporarily add any available unused resource (CPs, IFLs, ICFs, zIIPs, zAAPs, and SAPs) up to two times the purchased capacity. You can order this upgrade only using the CIU application through Resource Link.

The upgrade record is downloaded, staged, installed, and activated on your z10 BC server through its Support Element. The On/Off CoD record **is not** automatically activated when it is downloaded. It is placed in a "staged" area on the Support Element waiting to be installed and activated.

If you need the increased capacity for a longer period of time or you want to increase the capacity to the maximum amount order on the record, you can "replenish" the record. Using Resource Link, you place an order for a replenishment record to extend the expiration date, increase the capacity limits, or add additional tokens to an existing upgrade record. Replenishment allows you to update an existing record without having to place a completely new order and to update an existing record while capacity is active for that record.

When you order an On/Off CoD record, you can either prepay or post-pay for the upgrades. The payment method is based on the type of On/Off CoD upgrade you select:

- When you order a post-paid On/Off CoD record without spending limits, you select your upgrade configuration. You pay for what you activate during the activation time. You are charged on a 24-hour basis. For each month (starting with the month you activated the record), a report is generated. In the following month, you are billed for hardware and software charges.
- When you order a prepaid On/Off CoD record, you select your upgrade configuration and identify the duration of the configuration. Resource Link calculates the number of tokens you will need to activate your selected upgrade configurations. When the order is downloaded, you are billed for the total hardware cost. As resources are used, the corresponding number of tokens are decremented. Tokens are tracked on a 24-hour basis. For each month resources are used, a report is generated. In the following month, you are billed for software charges.
- When you order a post-paid On/Off CoD record with spending limits, you select your upgrade configuration and identify your spending limit for each upgrade.

Resource Link calculates the maximum number of tokens you may need to activate upgrade configurations without exceeding your spending limit. Tokens are tracked on a 24-hour basis. You will be notified when you are reaching the limit you set on your order. For each month (starting with the month you downloaded the record), a report is generated. In the following month, you are billed for hardware charges. Software charges are separate.

There are limits to the number of temporary zIIPs, zAAPs, IFLs, ICFs, and SAPs you can purchase. Refer to the *System z10 Capacity on Demand User's Guide* for details.

The On/Off CoD characteristics include:

- **Reusable On/Off CoD records** - Using a single On/Off CoD upgrade record, the z10 BC supports the moving from one capacity setting to another, either decreasing or increasing the amount of active temporary capacity. Multiple activations are possible within the same On/Off CoD upgrade record. The record remains valid for 180 days and may be replenished.
- **API used to activate** - z10 BC allows activation of On/Off CoD using the record identifier as an input parameter on the HMC SNMP API.
- **No-charge test** - The On/Off CoD test can be used to validate the processes to download, activate, and deactivate On/Off CoD capacity nondisruptively. With each On/Off CoD enabled machine, you are entitled to one no-charge test. The test may run for a maximum duration of 24 hours beginning with the activation of the test record. In addition to validating the On/Off CoD function, you can use this test as a training session for personnel who are authorized to activate On/Off CoD.
- **Multiple records simultaneously active** - An On/Off CoD record, CBU record, and CPE record can be active at the same time.
- **Administrative test** - For z10 BC, the Administrative Test function is no longer available.

Refer to the *System z10 Capacity on Demand User's Guide* for more information.

Capacity Backup (CBU)

Capacity Backup (CBU) is designed to replace lost capacity due to an emergency or disaster recovery situation. CBU increases capacity nondisruptively by allowing you to add specialty engines (IFLs, ICFs, zAAPs, zIIPs, SAPs) or add capacity by feature codes.

Each CBU record is allowed one 90-day "real" activation and a number of free 10-day test activations. The number of free test activations equates to the number of years that are purchased with the CBU record. (For example, a three year CBU record has three tests activations, a one year CBU record has one test activation.) Additional test activations beyond the free tests may be purchased in **single** increments up to a maximum of 15 CBU tests per record. This maximum of 15 tests per record cannot be exceeded and includes any free activations plus additional paid test activations.

The CBU characteristics include:

- **No password is required at time of activation.**
- **Specialty engines are managed by quantities** - Added capacity is dictated by processor types. You must indicate the number of engines that can be added to the permanent configuration.

- **CP capacity is managed by feature codes** - Feature codes either adds engines or increase the capacity to a permanent engine.
- **Choice in the length of contract** - Expiration date of a contract is 1 to 5 years. You have the capability to extend your CBU record up to the maximum 5 year limit. One test activation is provided for each additional CBU year added to the CBU record.
- **Limit on the number of zIIPs or zAAPs you can order** - This number cannot exceed the total number of permanents plus temporary CPs.

Refer to the *System z10 Capacity on Demand User's Guide* for more information.

Capacity for Planned Events (CPE)

Capacity for Planned Events (CPE) is designed to replace lost capacity for planned down time events, such as system migration or relocation (for a data center move). CPE increases capacity by allowing you to add model capacity or specialty engines (IFLs, ICFs, zAAPs, zIIPs, SAPs). Pricing is based on the model capacity and the type and quantity of the engines selected.

Each CPE order includes 1 activation for 3 days.

Refer to the *System z10 Capacity on Demand User's Guide* for more information.

Concurrent PU conversions

The z10 BC supports concurrent conversion of different processor unit (PU) types. This capability is extended to CPs, IFLs, ICFs, zIIPs, and zAAPs. This capability provides flexibility in configuring a z10 BC to meet the changing business environments.

Note: Concurrent PU conversion is not supported by CIU.

Reserved CP support in LPAR mode

With reserved CP support in LPAR mode, an LPAR may be defined with the number of logical CPs greater than the number of physical CPs. Additional CPs can be specified for the LPAR definition beyond the number of physical CPs currently installed on the model. Therefore, an enterprise planning to do a nondisruptive upgrade (with an LPAR defined of logical CPs equal to the number of physical CPs available on the installed hardware) does not need to deactivate, redefine, then reactivate in order to take advantage of the new CPs that have been activated. The enterprise simply needs to have defined additional CPs for the LPAR in advance. This ensures that any planned LPAR can be as large as the possible physical machine configuration, nondisruptively. For more information, refer to *System z10 Processor Resource/Systems Manager Planning Guide*.

Nondisruptive upgrades

The z10 BC Plan-Ahead process links the use of Capacity Upgrade on Demand with planning performed between IBM's account team and IBM's customer. Planning ahead enables customers to determine a future server configuration. IBM will also support its customers planning effort via capacity planning tools, IBM's order processing configurative and team sessions, with the objective of nondisruptive growth to satisfy essential capacity demand.

Processor capacity downgrades

You are allowed to downgrade your machine using CIU, CUoD, or MES. The primary benefit to downgrading is a reduction in software charges based on a lower reported machine capacity.

Some additional considerations should be noted when downgrading:

- Downgrades are done by “unassigning” either CPs or IFLs.
- There may be a charge to unassign and then reactivate engines.
- Unassigned engines are still owned by the customer.
- Unassigning unused engines can reduce software charges since many software products are priced based on the number of active engines.
- Unassigned engines can be reactivated by CIU, CUoD, or MES.
- Unassigned engines may be temporarily activated using On/Off CoD or CPE. When used as a temporary engine, unassigned engines can be used as any of the supported engine types (thus an unassigned IFL can be activated as a CP). Reduced hardware usage charges are available when using unassigned engines as the same type.
- Unassigning of engines and later reactivation is concurrent.

Appendix D. Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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

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

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

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

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

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol ([®] or [™]), these symbols indicate U.S. 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 "Copyright and trademark information" at www.ibm.com/legal/us/en/copytrade.shtml.

InfiniBand is a registered trademark of the InfiniBand Trade Association.

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

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.

Electronic emission notices

The following statement applies to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions contained in the installation manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, by installation or use of this equipment other than as specified in the installation manual, or by any other unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Canadian Department of Communications Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité aux normes du ministère des Communications du Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

European Union (EU) Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Warning: This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.

European Community contact:
IBM Technical Regulations
Pascalstr. 100, Stuttgart, Germany 70569
Telephone: 0049 (0) 711 785 1176
Fax: 0049 (0) 711 785 1283
email: tjahn@de.ibm.com

EC Declaration of Conformity (In German)

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:
"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraph 5 des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

Informationen in Hinsicht EMVG Paragraph 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

update: 2004/12/07

People's Republic of China Class A Compliance Statement

This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user may need to perform practical actions.

声 明

此为 A 级产品,在生活环境中,
该产品可能会造成无线电干扰。
在这种情况下,可能需要用户对其
干扰采取切实可行的措施。

Japan Class A Compliance Statement

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.

この装置は、クラスA情報技術装置です。この装置を家庭環境で使用する
と電波妨害を引き起こすことがあります。この場合には使用者が適切な
対策を講ずるよう要求されることがあります。 VCCI-A

Korean Class A Compliance Statement

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자
또는 사용자는 이점을 주의하시기 바라며, 만약 잘못 판매 또는
구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Taiwan Class A Compliance Statement

Warning: This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user will be required to take adequate measures.

警告使用者:
這是甲類的資訊產品,在
居住的環境中使用時,可
能會造成射頻干擾,在這
種情況下,使用者會被要
求採取某些適當的對策。

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Glossary

This glossary includes terms and definitions from:

- The *Dictionary of Computing*, SC20-1699.
- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology*, copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

The following cross-references are used in this glossary:

- **Contrast with.** This refers to a term that has an opposed or substantively different meaning.
- **See.** This refers the reader to multiple-word terms in which this term appears.
- **See also.** This refers the reader to terms that have a related, but not synonymous, meaning.
- **Synonym for.** This indicates that the term has the same meaning as a preferred term, which is defined in the glossary.

A

abend. Abnormal end of task

abnormal end of task (abend). Ending a task before its completion because of an error condition that cannot be resolved by recovery facilities while the task is being executed.

action. One of the defined tasks that an application performs. Actions modify the properties of an object or manipulate the object in some way.

activate logical partition. An operator-initiated procedure that performs a system reset to an LPAR and assigns the previously defined hardware to that partition. It causes an automatic IPL of the system control program to occur in the partition unless the operator performs the IPL manually. Contrast with *deactivate logical partition*.

active subchannel. A subchannel that is locked and either busy or has a pending interrupt, and is indicated by subchannel status word (SCSW) bit 24 equals 1. The control information resides in the channel subsystem because it is necessary for the current operation. Contrast with *inactive subchannel*. See also *busy subchannel*.

Note: An active subchannel can also reside in the local working storage of an IOP or channel.

active window. The window with which users are currently interacting. This is the window that receives keyboard input.

address. (1) A value that identifies a register, a particular part of storage, data source, or a data sink. The value is represented by one or more characters. (T) (2) To refer to a device or an item of data by its address. (I) (A) (3) The location in the storage of a computer where data is stored. (4) In data communication, the unique code assigned to each device or workstation connected to a network. (5) The identifier of a location, source, or destination.

address translation. (1) A value that identifies a register, a particular part of storage, a data source, or a data sink. The value is represented by one or more characters. (T) (2) In virtual storage systems, the process of changing the address of an item of data or an instruction from its virtual storage address to its real storage address. See also *dynamic address translation*.

ADMF. Asynchronous Data Mover Facility

alert. (1) A unit of information, usually indicating the loss of a system resource, passed from one machine or program to a host to signal an error. (2) An error message sent to the system services control point (SSCP) at the host system.

allocate. To assign a resource, such as a disk or a diskette file to perform a task. Contrast with *deallocate*.

American National Standard Code for Information Interchange (ASCII). The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity), used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphics characters. (A)

Note: IBM has defined an extension to ASCII code (characters 128 - 255).

ANSI. American National Standards Institute

APAR. Authorized program analysis report

API. Application programming interface

application. (1) The use to which an information processing system is put, for example, a payroll application, an airline reservation application, a network application. (2) A collection of software components used to perform specific types of work on a computer.

Application Assist Processor (AAP). A special processor configured for running Java applications on z10, z9, z990 and z890 class machines.

application program. (1) A program that is specific to the solution of an application problem. (T) (2) A program written for or by a user that applies to the user's work, such as a program that does inventory control or payroll. (3) A program used to connect and communicate with stations in a network, enabling users to perform application-oriented activities.

architected. Defined in the *z/Architecture Principles of Operation* or *IBM Enterprise Systems Architecture/390 Principles of Operation*.

ARP. Address Resolution Protocol

ASCII. American National Standard Code for Information Interchange

asynchronous. (1) Pertaining to two or more processes that do not depend upon the occurrence of specific events such as common timing signals. (T) (2) Without regular time relationship; unexpected or unpredictable with respect to the execution of program instructions. Contrast with *synchronous*.

ATM. Asynchronous transfer mode

ATM. Automatic teller machine

authorized program analysis report (APAR). A request for correction of a problem caused by a defect in a current release of a program unaltered by the user.

auto-answer. In data communication, the ability of a station to respond automatically to a call that it receives over a switched line.

auto-call. In data communication, the ability of a station to initiate a call automatically over a switched line.

B

basic mode. A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

batch. (1) An accumulation of data to be processed. (2) A group of records or data processing jobs brought together for processing or transmission. (3) Pertaining to activity involving little or no user action. Contrast with *interactive*.

BBU. Battery backup unit

BL. Parallel block multiplexer channel

block. A string of data elements recorded or transmitted as a unit. The element may be characters, words, or physical records. (T)

block multiplexer channel. A multiplexer channel that interleaves blocks of data. Contrast with *selector channel*. See also *byte multiplexer channel*

BPA. Bulk Power Assembly

buffer. (1) A routine or storage used to compensate for a difference in rate of flow of data, or time of occurrence of events, when transferring data from one device to another. (A) (2) To allocate and schedule the use of buffers. (A) (3) A portion of storage used to hold input or output data temporarily. See *fiber buffer*.

burst. In data communication, a sequence of signals counted as one unit in accordance with some specific criterion or measure. (A)

bus. (1) A facility for transferring data between several devices located between two end points, only one device being able to transmit at a given moment. (T) (2) A network configuration in which nodes are interconnected through a bidirectional transmission medium. (3) One or more conductors used for transmitting signals or power. (A)

byte. (1) A string that consists of a number of bits, treated as a unit, and representing a character. (T) (2) A binary character operated upon as a unit and usually shorter than a computer word. (A) (3) A string that consists of a particular number of bits, usually eight, that is treated as a unit, and that represents a character. (4) A group of eight adjacent binary digits that represent one extended binary-coded decimal interchange code (EBCDIC) character.

byte multiplexer channel. A multiplexer channel that interleaves bytes of data. Contrast with *selector channel*. See also *block multiplexer channel*.

C

CA. channel adapter (card)

cache. (1) A special purpose buffer storage, smaller and faster than main storage, used to hold a copy of the instructions and data obtained from main storage and likely to be needed next by the processor. (T) (2) A buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

cache structure. A coupling facility structure that enables high-performance sharing of cached data by multisystem applications in a Parallel Sysplex. Applications can use a cache structure to implement several different types of caching systems, including a store-through or a store-in cache.

CAW. Channel Address Word

CBU. Capacity Backup

CBY. ESCON byte multiplexer channel

CCC. Channel control check

CCW. Channel command word

CDC. Channel data check

central processor (CP). The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

central processor complex (CPC). The boundaries of a system, exclusive of I/O control units and devices, that can be controlled by a single operating system. A CPC consists of main storage, one or more central processor units, time-of-day clocks, and channels, which are or can be placed in a single configuration. A CPC also includes channel subsystems, service processors, and expanded storage where installed.

central storage. Storage that is an integral part of the processor and includes both main storage and the hardware system area.

CF. coupling facility (CF)

CFS. coupling facility sender

CFR. coupling facility receiver channel

CFRM. coupling facility resource management (policy)

CH. Channel card

channel. (1) A path along which signals can be sent, for example, input/output channel. (2) The system

element that controls one channel path, whose mode of operation depends on the type of hardware to which it is attached.

channel adapter. (1) A communication controller hardware unit used to attach the controller to a data channel. (2) Hardware that attaches a group of channels to the secondary data stager and prioritizes and stages data between the channels and the channel control element.

channel address. In S/370 mode, the 8 leftmost bits of an input/output address that identify the channel. See also *device address* and *input/output address*.

channel address word (CAW). An area in storage that specifies the location in main storage at which a channel program begins.

channel-attached. (1) Pertaining to attachment of devices directly by data channels (I/O channels) to a computer. (2) Pertaining to devices attached to a controlling unit by cables rather than by telecommunication lines. Contrast with *link-attached*.

channel command word (CCW). A doubleword at the location in main storage specified by the channel address word. One or more CCWs make up the channel program that directs data channel operations.

channel control check. A category of I/O errors affecting channel controls and sensed by the channel to which a device is attached. See also *channel data check*.

channel data check. A category of I/O errors, indicating a machine error in transferring data to or from storage and sensed by the channel to which a device is attached. See also *channel control check*.

channel data rate. The rate at which a channel can move data between a transmission link and processor storage during the data transfer portion of an I/O operation.

channel Licensed Internal Code. That part of the channel subsystem Licensed Internal Code used to start, maintain, and end all operations on the I/O interface. See also *IOP Licensed Internal Code*.

channel path (CHP). A single interface between a central processor and one or more control units along which signals and data can be sent to perform I/O requests.

channel path configuration. In an ESCON environment, the connection between a channel and a control unit or between a channel, an ESCON Director, and one or more control units. See also *link*, *point-to-point channel path configuration*, and *switched point-to-point channel path configuration*.

channel path identifier (CHPID). The channel subsystem communicates with I/O devices by means of a channel path between the channel subsystem and devices. A CHPID is a value assigned to each channel path of the System z that uniquely identifies that path. Up to 256 CHPIDs are supported for each channel subsystem

channel status word (CSW). An area in storage that provides information about the termination of input/output operations.

channel subsystem (CSS). A collection of subchannels that directs the flow of information between I/O devices and main storage, relieves the processor of communication tasks, and performs path management functions.

channel subsystem (CSS) Licensed Internal Code. Code that consists of the IOP Licensed Internal Code and the channel Licensed Internal Code.

channel-to-channel (CTC). Communication (transfer of data) between programs on opposite sides of a channel-to-channel adapter (CTCA).

channel-to-channel adapter (CTCA). An input/output device that is used by a program in one system to communicate with a program in another system.

check stop. The state that occurs when an error makes it impossible or undesirable to continue the operation in progress.

choice. An item that users may select. Choices appear in the selection fields, action bars, and pull-down menus.

CHP. Channel path

CHPID. See *channel path identifier*.

CI. Console integration

CIB. Coupling using InfiniBand

CICS. Customer Information Control System

CICS/ESA. Customer Information Control System/Enterprise Systems Architecture

CIU. Customer Initiated Upgrade

CKD. count key data.

click. To press and release a mouse button without moving the mouse pointer off the choice.

CLIST (command list). A data set in which commands and possibly subcommands and data are stored for subsequent execution.

CMOS. Complementary metal-oxide semiconductor

CNC. Mnemonic for an ESCON channel attached to an ESCON-capable device.

command. (1) A character string from a source external to a system that represents a request for system action. (2) A request from a terminal for performance of an operation or execution of a program. (3) A value sent on an I/O interface from a channel to a control unit that specifies the operation to be performed.

command chaining. The fetching of a new channel command word (CCW) immediately following the completion of the previous CCW.

command entry field. An entry field in which a user types commands.

command list. See *CLIST*.

command retry. A channel and control unit procedure that causes a command to be retried without requiring an I/O interrupt.

communication control unit. A communication device that controls transmission of data over lines in a network.

communication controller. (1) A device that directs the transmission of data over the data links of a network; its operation can be controlled by a program executed in a processor to which the controller is connected or it may be controlled by a program executed within the device. (T) (2) A type of communication control unit whose operations are controlled by one or more programs stored and executed in the unit. It manages the details of line control and the routing of data through a network.

complementary metal-oxide semiconductor (CMOS). A technology that combines the electrical properties of positive and negative voltage requirements to use considerably less power than other types of semiconductors.

concurrent maintenance. Hardware maintenance actions performed by a service representative while normal operations continue without interruption. See also *nondisruptive installation* and *nondisruptive removal*.

configuration. (1) The arrangement of a computer system or network as defined by the nature, number, and the chief characteristics of its functional units. More specifically, the term configuration may refer to a hardware configuration or a software configuration. (I) (A) (2) In an ESCON Director, the physical connection capability determined by a set of attributes. The attribute values specify the connectivity control status and identifiers associated with the ESCD and its ports. See also *active configuration*, *configuration matrix*, *connectivity attribute*, and *saved configuration*.

configure. To describe to the system the devices and optional features installed on the system.

connectivity. A term used to describe the physical interconnections of multiple devices/computers/networks employing similar or different technology or architecture together to accomplish effective communication between and among connected members involving data exchange or resource sharing.

console. A logical device used for communication between the user and the system. (A) See *display station, monitor console, operator console, program mode console, programming support console, service console, and system console.*

console integration (CI). The hardware and software facilities used to bring operating systems management and hardware systems management under a single control point.

control program. A computer program designed to schedule and to supervise the execution of programs of a computer system. (I) (A)

control unit. A hardware unit that controls the reading, writing, or displaying of data at one or more input/output units.

control unit data rate. The rate at which a control unit can move data between itself and a transmission link during the data transfer portion of an I/O operation.

controller. A unit that controls input/output operations for one or more devices.

conversational monitor system (CMS). A virtual machine operating system that provides general interactive time sharing, problem solving, and program development capabilities, and operates only under the VM control program.

Coordinated Server Time (CST). Represents the time in a CTN. Timekeeping messages carried over the coupling links determine the CST at each server.

Coordinated Timing Network (CTN). A collection of servers that are time synchronized to Coordinated Server Time (CST). All STP-configured servers in a CTN must have the same CTN ID.

coupling facility. A special partition that provides high-speed caching, list processing, and locking functions in a Parallel Sysplex.

coupling facility channel. A high bandwidth fiber optic channel that provides the high-speed connectivity required for data sharing between a coupling facility and the central processor complexes directly attached to it.

CP. (1) Control program (2) Central processor

CPC. Central processor complex

CPCID. CPC identifier

CPC image. Set of CPC resources that support a single control program.

CPU. Central processor unit

CPUID. CPU identifier

critical resource. A resource required for system operation. See also *system resource.*

CRW. Channel report word

CS. (1) Central storage. (2) Channel set. (3) Control storage.

CSS. Channel subsystem

CST. Coordinated Server Time

CSW. Channel status word

CTC. (1) Channel-to-channel. (2) Mnemonic for an ESCON or FICON channel attached to another ESCON or FICON channel respectively.

CTCA. Channel-to-channel adapter

CTN. Coordinated Timing Network

CU. Control unit

CUA. Control unit address.

CUADD. Control unit logical address.

CUoD. Capacity Upgrade on Demand

cursor. (1) A movable, visible mark used to indicate the position at which the next operation will occur on a display screen. (A) (2) A visual cue that shows the user where keyboard input will appear on the screen.

Customer Information Control System (CICS). An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

CVC. Mnemonic for an ESCON channel attached to a 9034.

D

DASD. See *direct access storage device.*

DASD subsystem. A storage control and its attached direct access storage devices.

DAT. Dynamic address translation

data processing (DP). The systematic performance of operations upon data; for example, arithmetic or logic

operations upon data, merging or sorting of data, assembling or compiling of programs. (T)

data rate. See *channel data rate*, *control unit data rate*, *device data rate*, *effective data rate*, and *input/output subsystem data rate*. See also *link rate*.

data sharing. The ability of concurrent subsystems (such as DB2 or IMS DB) or application programs to directly access and change the same data while maintaining data integrity.

data streaming. In an I/O interface, a mode of operation that provides a method of data transfer at up to 4.5 MB per second. Data streaming is not interlocked between the sender and the receiver. Once data transfer begins, the sender does not wait for acknowledgment from the receiver before sending the next byte. The control unit determines the data transfer rate.

data transfer. (1) The result of the transmission of data signals from any data source to a data receiver. (2) The movement, or copying, of data from one location and the storage of the data at another location.

data transfer mode. The method of information exchange used on an I/O interface. See *data streaming*.

DB2. DATABASE 2

DCA. Distributed Converter Assembly

DCAF. Distributed console access facility

DDR. Double Data Rate

deactivate logical partition. An operator-initiated procedure that releases the hardware assigned to a LPAR, making it available to other partitions. Contrast with *activate logical partition*.

Note: The operator should first deactivate the system control program, if possible or necessary, and then reactivate the partition, which could provide a reset to that partition, if required.

deallocate. To release a resource assigned to a task. Contrast with *allocate*.

DES. Data Encryption Standard

device. A mechanical, electrical, or electronic contrivance with a specific purpose.

device address. In S/370 mode, the 8 rightmost bits of an I/O address that identify a particular I/O device and a control unit on the designated channel. See *channel address*, *device-level addressing*, and *input/output address*.

DFSMS. Data Facility Storage Management Subsystem.

direct access storage. A storage device that provides direct access to data. (I) (A) See also *random access memory*.

direct access storage device (DASD). (1) A storage device in which the location of each data record can be directly addressed. (2) A device in which the access time is effectively independent of the location of the data. (Restriction: Does not refer to diskette drive.)

DP. data processing.

DSS. Digital Signature Standard

dual inline memory module (DIMM). A small circuit board with memory-integrated circuits containing signal and power pins on both sides of the board.

dynamic address translation (DAT). In virtual storage systems, the change of a virtual storage address to a real storage address during execution of an instruction. See also *address translation*.

dynamic reconfiguration management. In MVS™, the ability to modify the I/O configuration definition without needing to perform a power-on reset (POR) of the hardware or an initial program load (IPL).

dynamic storage reconfiguration. A PR/SM LPAR function that allows central or expanded storage to be added or removed from an LPAR without disrupting the system control program operating in the LPAR.

E

EC. Engineering change

ECC. Error checking and correction

ECKD. Extended count key data

EEPROM. Electrically erasable programmable read only memory

EIA. Electronics Industries Association. One EIA unit is 1.75 inches or 44.45mm.

Enterprise Systems Connection (ESCON). A set of products and services that provides a dynamically connected environment using optical cables as a transmission medium.

EPO. Emergency power off

error checking and correction (ECC). In a processor, the detection and correction of all single-bit errors, plus the detection of double-bit and some multiple-bit errors.

ESA. (1) Enterprise Systems Architecture (2) Expanded storage array

ESA/370. Enterprise Systems Architecture/370

ESA/390. Enterprise Systems Architecture/390

ESCD. Enterprise Systems Connection (ESCON) Director

ESCM. Enterprise Systems Connection Manager

ESCON. Enterprise Systems Connection

ESCON channel. A channel having an Enterprise Systems Connection channel-to-control-unit I/O interface that uses optical cables as a transmission medium. Contrast with *parallel channel*.

ESCON Director (ESCD). A device that provides connectivity capability and control for attaching any two links to each other.

ESCON environment. The data processing environment having an Enterprise Systems Connection channel-to-control-unit I/O interface that uses optical cables as a transmission medium.

ESCON Manager (ESCM). A licensed program that provides host control and intersystem communication capability for ESCON Director connectivity operations.

Ethernet definition. A communication network (USA, Xerox 1975).

ETR. External Time Reference

event. (1) An occurrence or happening. (2) An occurrence of significance to a task; for example, the completion of an asynchronous operation, such as an input/output operation.

expanded storage. Optional high-speed storage that transfers 4 KB pages to and from central storage.

F

facility. (1) An operational capability, or the means for providing such a capability. (T) (2) A service provided by an operating system for a particular purpose; for example, the checkpoint/restart facility.

FCP. Fibre Channel Protocol for SCSI

FDDI. Fiber Distributed Data Interface (100Mbps/second fiber optic LAN)

fiber. See *optical fiber*.

FICON. Fibre CONnection

frame. (1) For a zSystem or zSeries microprocessor cluster, a frame contains one or two central processor complexes (CPCs), Support Elements, and AC power distribution. (2) A housing for machine elements. (3) The hardware support structure, covers, and all electrical parts mounted there in that are packaged as one entity for shipping. (4) A formatted display. See *display frame* and *transmission frame*.

FRU. Field-replaceable unit.

ft. Foot.

G

Gb. Gigabit

GB. Gigabyte.

GbE. Gigabit Ethernet.

gigabit (Gb). A unit of measure for storage size. One gigabit equals one billion bits.

Gigabit Ethernet. An OSA channel (CHPID type OSD)

gigabyte (GB). (1) A unit of measure for storage size. One gigabyte equals 1,073,741,824 bytes. (2) Loosely, one billion bytes.

GMT. Greenwich mean time.

H

Hardware Management Console. A user interface through which data center personnel configure, control, monitor, and manage System z hardware and software resources. The HMC communicates with each central processor complex (CPC) through the CPC's Support Element (SE).

hardware system area (HSA). A logical area of central storage, not addressable by application programs, used to store Licensed Internal Code and control information.

HCA. Host Channel Adapter

HCA1-O fanout. The HCA1-O (optical) fanout card is used for coupling using an InfiniBand connection on a z9. The HCA1-O fanout is designed to support a two-port 12x IB-SDR optical link operating at a link rate of 3 Gbps.

HCA2-C fanout. The HCA2-C (copper) fanout card has InfiniBand connections used for internal I/O on a z10. The HCA2-C fanout is designed to support a two-port 12x IB-DDR copper link operating at a link rate of 6 Gbps.

HCA2-O fanout. The HCA2-O (optical) fanout card is used for coupling using an InfiniBand connection on a z10 server. The HCA2-O fanout is designed to support a two-port 12x IB-DDR optical link operating at a link rate of 6 Gbps.

HCD. Hardware configuration definition

HDD. HDD

| **HiperSockets network traffic analyzer (HS NTA).**
 | Trace HiperSockets network traffic to help simplify
 | problem isolation and resolution. Supported on System
 | z10.

HSA. Hardware system area

HMCA. Hardware Management Console Application

|

IBB. Internal Bus Buffer

IBF. Internal Battery Feature

IC. Internal Coupling link

ICB. Integrated Cluster Bus link

ICF. Internal Coupling Facility

ICMF. Integrated Coupling Migration Facility

ICRF. Integrated Cryptographic Facility

ICSF. Integrated Cryptographic Service Facility

IDAW. Indirect data address word

IFB. InfiniBand

IFB-MP (InfiniBand Multiplexer). On a System z10, used for an I/O cage intraconnection to I/O features.

IFCC. Interface control check

IFL. Integrated Facility for Linux

IML. Initial machine load

IMS. Information Management System

initial machine load (IML). A procedure that prepares a device for use.

initial program load (IPL). (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage at the beginning of a work day or after a system malfunction. (3) The process of loading system programs and preparing a system to run jobs.

initialization. (1) The operations required for setting a device to a starting state, before the use of a data medium, or before implementation of a process. (T) (2) Preparation of a system, device, or program for operation. (3) To set counters, switches, addresses, latches, or storage contents to zero or to other starting values at the beginning of, or at the prescribed points in, a computer program or process.

Integrated Facility for Applications (IFA). A general purpose assist processor for running specific types of applications. See Application Assist Processor (AAP).

input/output (I/O). (1) Pertaining to a device whose parts can perform an input process and an output process at the same time. (I) (2) Pertaining to a functional unit or channel involved in an input process, output process, or both, concurrently or not, and to the data involved in such a process. (3) Pertaining to input, output, or both.

input/output configuration. The collection of channel paths, control units, and I/O devices that attach to the processor complex.

input/output configuration data set (IOCDS). The data set that contains an I/O configuration definition built by the I/O configuration program (IOCP).

input/output configuration program (IOCP). A program that defines to a system all the available I/O devices and the channel paths.

input/output configuration source. The file used as input to the IOCP when an IOCDS is created. It is an 80-column card-image, and is initially provided in EBCDIC form on a system tape, or in ASCII form on a 3.5-inch diskette.

interrupt. (1) A suspension of a process, such as execution of a computer program caused by an external event, and performed in such a way that the process can be resumed. (A) (2) To stop a process in such a way that it can be resumed. (3) In data communication, to take an action at a receiving station that causes the sending station to end a transmission. (4) To temporarily stop a process.

I/O. See *input/output*.

IOCDS. I/O configuration data set

IOCP. I/O configuration program

IODF. I/O definition file

IPL. See *initial program load*.

IPv6. Internet Protocol Version 6

ISC. InterSystem Channel

ISDN. Integrated-Services Digital Network

K

KB. Kilobyte

kilobyte. (1) A unit of measure for storage size. (2) Loosely, one thousand bytes.

km. Kilometer

L

LAN. See *local area network*.

laser. A device that produces optical radiation using a population inversion to provide light amplification by stimulated emission of radiation and (generally) an optical resonant cavity to provide positive feedback. Laser radiation can be highly coherent temporally, or spatially, or both. (E)

LCSS. Logical channel subsystem

LED. Light-emitting diode

LIC. Licensed Internal Code

Licensed Internal Code (LIC). Software provided for use on specific IBM machines and licensed to customers under the terms of IBM's Customer Agreement.

light-emitting diode (LED). A semiconductor chip that gives off visible or infrared light when activated.

local area network (LAN). A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary can be subject to some form of regulation.

Note: A LAN does not use store and forward techniques.

logical address. The address found in the instruction address portion of the program status word (PSW). If translation is off, the logical address is the real address. If translation is on, the logical address is the virtual address. See also *absolute address*, *physical address*, *real address*, and *virtual address*.

logical control unit. A group of contiguous words in the hardware system area that provides all of the information necessary to control I/O operations through a group of paths that are defined in the IOCDS. Logical control units represent to the channel subsystem a set of control units that attach common I/O devices.

logical partition (LPAR). A subset of the processor hardware that is defined to support the operation of a system control program (SCP).

logical processor. In LPAR mode, central processor resources defined to operate in an LPAR like a physical central processor.

logical unit (LU). In SNA, a port to the network through which an end user accesses the SNA network and the functions provided by system services control points (SSCPs). An LU can support at least two

sessions - one with an SSCP and one with another LU - and may be capable of supporting many sessions with other LUs.

logically partitioned (LPAR) mode. A central processor complex (CPC) power-on reset mode that enables use of the PR/SM feature and allows an operator to allocate CPC hardware resources (including central processors, central storage, expanded storage, and channel paths) among LPARs. Contrast with *basic mode*.

LU. Logical unit

LUPS. Local uninterruptible power supply

M

MAC. Message Authentication Code

main storage. (1) Program-addressable storage from which instructions and other data can be loaded directly into registers for subsequent processing. (I) (A) (2) That part of internal storage into which instructions and other data must be loaded for subsequent execution or processing. (3) The part of a processor unit where programs are run. See *central storage*.

Notes:

1. Main storage refers to the whole program-addressable execution space and can include one or more storage devices.
2. The term *main storage* is generally used in large and intermediate computers. The term *memory* is primarily used in microcomputers, calculators, and some minicomputers.

maintenance change level (MCL). A change to correct a single licensed internal code design defect. Higher quality than a *patch*, and intended for broad distribution. Considered functionally equivalent to a software PTF.

MAU. Multistation access unit

Mb. Megabit

MB. Megabyte

MBA. Memory bus adapter

MBA fanout. On a z9, an MBA fanout card is used for coupling using ICB-4 and for connecting an I/O cage to a CPC cage. On a z10, MBA fanout is used for coupling using ICB-4.

MCCU. Multisystem channel communication unit

MCL. See *maintenance change level*.

megabit (Mb). A unit of measure for storage size. One megabit equals 1,000,000 bits.

megabyte (MB). (1) A unit of measure for storage size. One megabyte equals 1,048,576 bytes. (2) Loosely, one million bytes.

menu bar. The area at the top of the primary window that contains keywords that give users access to actions available in that window. After users select a choice in the action bar, a pulldown menu appears from the action bar.

MIDAW. Modified Data Indirect Address Word

MIF. Multiple Image Facility

modem (modulator/demodulator). A device that converts digital data from a computer to an analog signal that can be transmitted on a telecommunication line, and converts the analog signal received to data for the computer.

mouse. In computer graphics, a pointing device operated by moving it on a surface.

Multiple Image Facility (MIF). A facility that allows channels to be shared among PR/SM LPARs in an ESCON or FICON environment.

| **multichip module (MCM).** The fundamental processor
| building block for System z. Each System z “book” is
| comprised of a glass ceramic multichip module of
| processor units (PUs) and memory cards, including
| multilevel cache memory.

multiplexer channel. A channel designed to operate with a number of I/O devices simultaneously. Several I/O devices can transfer records at the same time by interleaving items of data. See *block multiplexer channel* and the *multiplexer channel*.

MVS. Multiple Virtual Storage

MVS image. A single occurrence of the MVS/ESA™ operating system that has the ability to process work.

MVS system. An MVS image together with its associated hardware, which collectively are often referred to simply as a system, or MVS system.

N

NetBIOS. Local area network basic input/output system

network. (1) An arrangement of nodes and connecting branches. (T) (2) A configuration of data processing devices and software connected for information exchange.

NIC. Numerically intensive computing

O

On/Off Capacity on Demand (On/Off CoD). Used to temporarily turn on CPs, IFLs, ICFs, zIIPs, and zAAPs.

operate. To do a defined action, such as adding or comparing, performed on one or more data items.

operating system (OS). Software that controls the execution of programs and that may provide services such as resource allocation, scheduling, input/output control, and data management. Although operating systems are predominantly software, partial hardware implementations are possible. (T)

operator console. (1) A functional unit containing devices that are used for communications between a computer operator and a computer. (T) (2) A display used for communication between the operator and the system, used primarily to specify information concerning application programs and I/O operations and to monitor system operation.

optical cable. A fiber, multiple fibers, or a fiber bundle in a structure built to meet optical, mechanical, and environmental specifications. (E) See also *jumper cable*, *optical cable assembly*, and *trunk cable*.

optical fiber. Any filament made of dielectric materials that guides light, regardless of its ability to send signals. (E) See also *fiber optics* and *optical waveguide*.

OS. Operating system

OSA. Open Systems Adapter (OSA-Express3 and OSA-Express2). The OSA is an integrated hardware feature that provides direct connection to clients on local area networks (LANs).

OSA/SF. Open Systems Adapter/Support Facility

P

parallel channel. (1) A channel having a S/360™ and S/370 channel-to-control-unit I/O interface that uses bus-and-tag cables as a transmission medium. Contrast with *ESCON channel*. (2) A data path along which a group of signals representing a character or any other entity of data can be sent simultaneously.

Parallel Sysplex. A set of MVS systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads. See also *MVS system*.

PIN. Personal Identification Number

PKA. Public-Key-Algorithm

PKSC. Public-Key Secure Cable

point-to-point channel path configuration. In an I/O interface, a configuration that consists of a single link between a channel and one control unit. Contrast with *switched point-to-point channel path configuration*.

point-to-point connection. A connection established between two data stations for data transmission. Contrast with *multipoint connection*.

Note: The connection may include switching facilities.

POR. Power-on reset

power-on reset. A function that re-initializes all the hardware in the system and loads the internal code that enables the machine to load and run an operating system. This function is intended as a recovery function.

power-on reset state. The condition after a machine power-on sequence and before an IPL of the control program.

problem analysis. The process of identifying the source of a problem; for example, a program component, a machine failure, telecommunication facilities, user or contractor-installed programs or equipment, an environment failure such as a power loss, or a user error.

problem determination (PD). The process of determining the source of a problem; for example, a program component, machine failure, telecommunication facilities, user or contractor-installed programs or equipment, environmental failure such as a power loss, or user error.

problem management. The management discipline that handles a problem from its detection through its final resolution. Problem management is composed of the following:

- Problem determination
- Problem diagnosis
- Problem bypass and recovery
- Problem resolution
- Problem tracking and control.

processor. (1) In a computer, a functional unit that interprets and executes instructions. A processor consists of at least an instruction control unit and an arithmetic and logic unit. (T) (2) The functional unit that interprets and executes instructions. (3) The boundaries of a system, exclusive of I/O control units and devices, that can be controlled by a single operating system. A processor consists of main storage, one or more central processors, time-of-day clocks, and channels, which are, or can be, placed in a single configuration. A processor also includes channel subsystems, and expanded storage where installed.

processor complex. A system configuration that consists of all the machines required for operation; for

example, a processor unit, a processor controller, a system display, a service support display, and a power and coolant distribution unit.

processor console. The workstation from which an operator can control and observe system hardware operation. See also system console.

Processor Resource/Systems Manager (PR/SM) facility. The feature that allows the processor to use several system control programs (SCPs) simultaneously, provides logical partitioning capability for the real machine, and provides support for multiple preferred guests.

processor unit (PU). A PU can be defined as a CP, ICF, IFL, zIIP, zAAP or spare SAP.

profile. (1) A description of the characteristics of an entity to which access is controlled. (2) Data that describes the significant characteristics of a user, a group of users, or one or more computer resources.

program. Sequence of instructions for a computer. A program interacts and relies on either the hardware or other programs.

program status word (PSW). An area in storage used to indicate the sequence in which instructions are executed, and to hold and indicate the status of the computer system.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM as resulting from a defect in a current, unaltered release of the program.

PR/SM. Processor Resource/Systems Manager

PSC. Power Sequence Controller

PSP. Preventive service planning

PSW. Program status word

PTF. Program temporary fix

R

RAS. Reliability, availability, serviceability

reconfiguration. (1) A change made to a given configuration in a computer system; for example, isolating and bypassing a defective functional unit or connecting two functional units by an alternative path. Reconfiguration is effected automatically or manually and can be used to maintain system integrity. (T) (2) The process of placing a processor unit, main storage, and channels offline for maintenance, and adding or removing components. (3) Contrast with *nondisruptive installation* and *nondisruptive removal*.

recovery. To maintain or regain system operation after a failure occurs. Generally, to recover from a failure is to identify the failed hardware, to deconfigure the failed hardware, and to continue or restart processing.

remote service facility (RSF). (1) A control program plus associated communication equipment that allows local personnel to connect to an IBM service center, and allows remote personnel to operate the remote system or send new internal code fixes to it, if properly authorized. (2) A system facility invoked by Licensed Internal Code that provides procedures for problem determination and error detection.

Remote Technical Assistance and Information Network (RETAIN). A database, accessible to service representatives, of information relating to IBM-installed products.

RETAIN. Remote Technical Assistance and Information Network

REXX. Restructured extended executor language

ring network. A network configuration in which devices are connected by unidirectional transmission links to form a closed path.

Note: A ring of an IBM token-ring network is referred to as a LAN segment or as a token-ring network segment.

RMF. Resource Measurement Facility™

RPQ. Request for Price Quotation

RPS. Rotational positional sensing/sensor

RSA. Rivest-Shamir-Adelman

RSF. Remote support facility

S

SAD. System Activity Display

SAP. System Assist Processor

SCP. System control program

scroll. To move a display image vertically or horizontally to display data that cannot be observed within the boundaries of the display frame.

scroll bar. A window component associated with a scrollable area that provides users a visual cue that more information is available and that the unseen information can be manipulated into view using the mouse. Users scroll the information in the window by interacting with the scroll bar.

SCSI. Small Computer System Interface

SDLC. See *synchronous data link control*.

SDR. Single Data Rate

SE. Support Element

Server Time Protocol (STP). A message based protocol designed to enable multiple servers to maintain time synchronization with each other. The timekeeping information is passed over data links (externally defined coupling links) between servers. It provides time synchronization for the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers and CFs without requiring the Sysplex Timer.

service representative. A person who performs maintenance services for IBM hardware products or systems. See also *IBM program support representative*.

SIE. Start Interpretive Execution

single point of control. The characteristic a Parallel Sysplex displays when you can accomplish a given set of tasks from a single workstation, even if you need multiple IBM and vendor products to accomplish that particular set of tasks.

single system image. The characteristic a product displays when multiple images of the product can be viewed and managed as one image.

SNA. See systems network architecture.

SNA/Management Services (SNA/MS). Functions distributed among network components to operate, manage, and control the network.

SNA/MS. SNA/Management Services.

SNA network. The part of a user-application network that conforms to the formats and protocols of Systems Network Architecture. It enables reliable transfer of data among end-users and provides protocols for controlling the resources of various network configurations. The SNA network consists of network addressable units (NAUs), boundary function components, and the path control network.

SNMP. Simple network management protocol

STI. Self-Timed Interconnect

STI-MP (Self-Timed Interconnect Multiplexer). For System z9, used for an I/O cage intraconnection.

STP. Server Time Protocol

storage. (1) A functional unit into which data can be placed, in which they can be retained, and from which they can be retrieved. (T) (2) The action of placing data into a storage device. (I) (A)

structure. A construct used by MVS to map and manage storage on a coupling facility. See cache structure, list structure, and lock structure.

subchannel. In 370-XA, ESA/390 modes, and z/Architecture modes, the facility that provides all of the information necessary to start, control, and complete an I/O operation.

subchannel number. A system-unique 16-bit value used to address a subchannel. See also *channel path identifier*, *device identifier*, and *device number*.

subsystem. A secondary or subordinate system, or programming support, usually capable of operating independently of or asynchronously with a controlling system. (T) See *DASD subsystem*. and *storage subsystem*.

subsystem storage. Synonym for *cache*.

Support Element (SE). (1) An internal control element of a processor that assists in many of the processor operational functions. (2) A hardware unit that provides communications, monitoring, and diagnostic functions to a central processor complex (CPC).

synchronous data link control (SDLC). A form of communication line control that uses commands to control data transfer over a communication line.

Sysplex Timer. An IBM unit that synchronizes the time-of-day (TOD) clocks in multiple processors or processor sides. External Time Reference (ETR) is the MVS generic name for the IBM Sysplex Timer.

system. Comprises the processor complex and all attached and configured I/O and communication devices.

system area. A logical area of central storage used to store Licensed Internal Code and control information (not addressable by application programs).

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operational sequences for transmitting information units through, and controlling the configuration and operation of, networks.

S/370. IBM System/370

S/390. IBM System/390®

T

target logical partition. In LPAR mode, the target logical partition is the LPAR that is selected on the System Monitor panel. It is the current or immediate LPAR; the LPAR that the operator is working with or looking at.

target processor. The processor that controls execution during a program restart, instruction trace, stand-alone dump, or IPL, and whose ID is identified by highlighting on the status line.

TCP/IP. Transmission Control Protocol/Internet Protocol

TDES. Triple Data Encryption Standard

time-of-day (TOD) clock. A system hardware feature that is incremented once every microsecond, and provides a consistent measure of elapsed time suitable for indicating date and time. The TOD clock runs regardless of whether the processor is in a running, wait, or stopped state.

TKE. Trusted Key Entry

TOD. Time of day

token. A sequence of bits passed from one device to another on the token-ring network that signifies permission to transmit over the network. It consists of a starting delimiter, an access control field, and an end delimiter. The access control field contains a bit that indicates to a receiving device that the token is ready to accept information. If a device has data to send along the network, it appends the data to the token. When data is appended, the token then becomes a frame.

token-ring network. (1) A ring network that allows unidirectional data transmission between data stations, by a token passing procedure, such that the transmitted data return to the transmitting station. (T) (2) A network that uses ring topology, in which tokens are passed in a circuit from node to node. A node that is ready to send can capture the token and insert data for transmission.

Note: The IBM token-ring network is a baseband LAN with a star-wired ring topology that passes tokens from network adapter to network adapter.

TPF. Transaction processing facility

transaction processing. In batch or remote batch processing, the processing of a job or job step. In interactive processing, an exchange between a terminal and another device that does a particular action; for example, the entry of a customer's deposit and the updating of the customer's balance.

TSCF. Target system control facility

U

UCW. Unit control word

UPC. Universal power controller

UPS. Uninterruptible power supply

user interface. Hardware, software, or both that allows a user to interact with and perform operations on a system, program, or device.

V

VLAN. Virtual Local Area Network

VSE. Virtual Storage Extended

W

WAC. Wide Area Connector (card)

window. (1) An area of the screen with visible boundaries through which information is displayed. A window can be smaller than or equal in size to the screen. Windows can overlap on the screen and give the appearance of one window being on top of another. (2) A choice in the action bar of some applications. Users select it to arrange the display of several windows or to change the active window. (3) A choice in the action bar of multiple-document interface applications. (4) A choice in an action bar that allows the user to arrange the display of all open windows and to change the active window. (5) A choice in an action bar of multiple-document interface applications that allows a user to arrange the display of all open windows and to change the active window.

workstation. (1) A functional unit at which a user works. A workstation often has some processing capability. (T) (2) A terminal or microcomputer, usually one that is connected to a mainframe or network, at which a user can perform applications. (3) See also *display station* and *terminal*.

Z

zAAP. (1) System z Application Assist Processor, System z10 Application Assist Processor (2) A specialized processor that provides a Java execution environment, which enables Java-based web applications to be integrated with core z/OS business publications and backend database systems.

zIIP. (1) System z Integrated Information Processor, System z10 Integrated Information Processor. (2) A specialized processor that provides computing capacity for selected data and transaction processing workloads, and for selected network encryption workloads.

z10 BC. IBM System z10 Business Class server

z10 EC. IBM System z10 Enterprise Class server

z800. eServer zSeries 800

z890. eServer zSeries 890

z900. eServer zSeries 900

z9 BC. IBM System z9 Business Class server

z9 EC. IBM System z9 Enterprise Class server

Numerics

370-XA. IBM System/370 extended architecture

Index

Numerics

3390 61
64-bit addressing 82

A

A frame 11
activation 113
adapters
 maximums 31
addressing
 FCP 51
 network concentrator 71
AID
 assignments 38
API (Application Programming Interfaces) 112
Application Programming Interfaces (API) 112
architecture
 ESA/390 8
 z/Architecture 8
assignments
 AID 38
 CHPID 37, 38
 PCHID 37, 38
ATM 98
availability guidelines 34

B

Backup Time Server 126
block multiplexer mode 55
BPA (Bulk Power Assembly) 21
BPC 21
BPE (Bulk Power Enclosure) 21
BPF (Bulk Power Fan) 21
BPH (Bulk Power Hub) 21
BPI (Bulk Power Interface) 21
BPR (Bulk Power Regulator) 21
Bulk Power Assembly (BPA) 21
Bulk Power Enclosure (BPE) 21
Bulk Power Fan (BPF) 21
Bulk Power Hub (BPH) 21
Bulk Power Interface (BPI) 21
Bulk Power Regulator (BPR) 21
burst mode 55
byte multiplexer
 mode 55
 operation 55

C

cable ordering 101
cabling
 fiber optic 7
 report 102
 responsibilities 100
cache-hit rates 63

Cancel I/O (XSCH) instruction 57
Capacity Backup (CBU) 128, 139
capacity upgrades 137
cards
 DCA 18, 20
 ISC-D 77
 ISC-M 77
 memory 16
 OSC/ETR 17
 PSC24V 20
cascaded directors 125
CBU (Capacity Backup) 128
CCC (Channel Control Check) 58
CCW (Channel Command Word) 57
CDC (Channel Data Check) 58
Central Processor (CP) 13, 23
Central Processor Complex (CPC) 12
central storage 15
certification
 EAL5 91
 FIPS 98
CF (coupling facility) 80
CF duplexing 82, 86
CFCC (Coupling Facility Control Code) 80
 48 internal tasks 82
 64-bit addressing 82
 CF duplexing 82
 considerations 81
 Level 11 83
 Level 12 82
 Level 13 82
 Level 14 82
 Level 16 82
 patch apply 81
chaining operations 57
channel
 commands 56
 path definition 42
 time out 59
Channel Command Word (CCW) 57
Channel Control Check (CCC) 58
Channel Data Check (CDC) 58
channel hardware
 FCP 50
channel program characteristics 63
channel sharing
 FCP 52
channel subsystem
 planning 37
channel subsystem (CSS) 22, 31, 32
channel subsystem workload characteristics 63
channel-to-channel connection 59
channels 23
 dedicated 23
 ESCON 54
 HiperSockets 41
 internal coupling (IC) 41
 InterSystem Coupling-3 (ISC-3) 77

- channels (*continued*)
 - IOCP definitions 32
 - maximums 31
 - peer 33
 - reconfigurable 23
 - shared 23, 24
 - spanned 23, 41
- characteristics
 - cache hit rates 63
 - channel program 63
 - channel subsystem workload 63
 - I/O device I/O 63
- CHPID
 - assignments 37, 38
 - types 32
- CHPID Mapping Tool 40
- CIU (Customer Initiated Upgrade) application 137
- classes
 - device 56
- Clear Subchannel (CSCH) instruction 57
- clustering 89
- command
 - chaining 58
- comments
 - sending xviii
- compatibility
 - programming 27
- concurrent
 - channel upgrade 125
 - maintenance
 - channels 127
 - connectivity links 127
 - cryptographic 127
 - Hardware Management Console/SE 127
 - OSA 127
 - power 127
 - undo CBU 129
- configurations
 - ICB-4 36
 - InfiniBand 36
 - system 11
- connection
 - CTC 59
- connectivity
 - subchannel 34
- considerations 83
 - coupling connection 83
 - HCD 43
 - IOCP 42
- control check operation 58
- control unit priority 60
- converter operation
 - ESCON 54
- coupling facility (CF) 80
 - duplexing 86
- Coupling Facility Control Code (CFCC) 80, 81
- coupling link
 - peer channels 33
- CP (Central Processor) 13, 23
- CP Assist for Cryptographic Function (CPACF) 93
- CPACF (CP Assist for Cryptographic Function) 93

- CPC (Central Processor Complex) 12
- CPC drawer
 - PCHID 38
- critical time 56
- Crypto Express2 95
- cryptography 93
 - CPACF 93
- CSS (channel subsystem) 22, 31, 32
- CTC connection 59
- CUoD (Capacity Upgrade on Demand)
 - Reserved CP support 140
- Customer Initiated Upgrade (CIU) application 137
- customization 108

D

- data
 - chaining 57
 - streaming protocol 55
- data check operation 58
- data rates
 - channel 62
 - control unit 62
 - device 62
 - I/O subsystem 62
- data transfer rate 62
- DCA (Distributed Converter Assembly) cards 12, 18, 20
- dedicated channels 23
- degrade indicator 109
- device
 - class 56
 - I/O 51, 55
 - performance 55
 - sharing 24
- device sharing 23
 - FCP 52
- Distributed Converter Assembly (DCA) cards 12, 18, 20
- drawer
 - positions 12
- drawers
 - I/O 18
- Dynamic Channel Path Management 90
- dynamic I/O configuration 125
- dynamic link aggregation 69
- dynamic reconnection 61
- dynamic storage reconfiguration 23

E

- ECC (Error Checking and Correction) 16, 124
- ending status 58
- Error Checking and Correction (ECC) 16
- Error Correction Code (ECC) 124
- error handling 58
- ESA/390 mode I/O instructions 57
- ESCON
 - channel commands 56
 - channel multiplexing modes 55
 - channels 54

ESCON (*continued*)
 converter operation 54
 I/O operations control 56
ESCON channels
 performance characteristics 63
ethernet switch 21
ETR (External Time Reference) 126
ETR (External Time Reference) card 17
ETR card 12
expanded storage 15
External Time Reference (ETR) 126
External Time Reference (ETR) card 17

F

fanout cards
 HCA 17
 MBA 17
FCP (Fibre Channel Protocol)
 addressing 51
 channel hardware 50
 channel sharing 52
 device sharing 52
 for SCSI 49
 positioning 52
features
 I/O 18
fiber optic cabling 7
Fiber Quick Connect (FQC) 100
Fibre Channel Protocol (FCP)
 for SCSI 49
FICON
 cascaded directors 125
FICON Express 47
FICON Express2 47
FICON Express4 46
FICON Express8 45
FQC (Fiber Quick Connect) 100
frame, A 11
frames, system 11
FSP card 12
functions
 Parallel Sysplex 75

G

GDPS (Geographically Dispersed Parallel Sysplex) 87, 129
Geographically Dispersed Parallel Sysplex (GDPS) 87, 129

H

Halt Subchannel (HSCH) instruction 57
Hardware Configuration Definition (HCD) 43, 111
Hardware Management Console 25, 105
 availability 123
 capabilities 105
 function 105
Hardware Management Console and Support Element
 wiring options 108

Hardware Management Console Application
 (HWMCA) 106
hardware messages 109
Hardware System Area (HSA) 16
HCA (host channel adapter) fanout cards 17
HCD (Hardware Configuration Definition) 43, 111
highlights 2
HyperSockets
 CHPID 41
 I/O connectivity 70
 network traffic analyzer (NTA) 71, 131
host channel adapter (HCA) fanout cards 17
HSA (Hardware System Area) 16
HWMCA (Hardware Management Console
 Application) 106

I

I/O
 device 51, 55
 device definition 43
 drawers 18, 35
 features 18
 interface mode 55
 interface protocol 59
 interruptions
 CCC 58
 CDC 58
 IFCC 58
 operations control 56
 PCHID 38
 performance 55
 system reset 58
I/O configuration 83
I/O device characteristics 63
I/O priority queuing (IOPQ) 89
I/O Subsystem (IOSS) 32
IBF (Internal Battery Feature) 21
IBM Resource Link 135
ICB-4 (Integrated Cluster Bus-4) 36, 77, 78
ICF (Internal Coupling Facility) 3, 13, 23
IFB (InfiniBand) 36
IFCC (Interface Control Check) 58
IFL (Integrated Facility for Linux) 2, 23
InfiniBand (IFB) 36
instructions
 SIE (start interpretive execution) 28
Integrated Cluster Bus-4 (ICB-4) 36, 78
Integrated Facility for Linux (IFL) 23
Intelligent Resource Director (IRD) 88
interface
 protocol, I/O 55
Interface Control Check (IFCC) 58
interface mode
 data streaming 55
 interlocked 55
interlocked
 protocol 55
Internal Battery Feature (IBF) 21
internal coupling (IC)
 channels 41

- internal coupling (IC) (*continued*)
 - links 80
- Internal Coupling Facility (ICF) 13, 23
- interruptions
 - control check 58
 - data check 58
 - I/O 58
 - machine check 58
- InterSystem Coupling-3 (ISC-3) channels 77
- IOCP
 - channel definitions 32
 - considerations 42
- IOSS (I/O Subsystem) 32
- iQDIO (internal Queued Direct Input/Output) 70, 72
- IRD (Intelligent Resource Director) 88
- ISC-3 (InterSystem Coupling-3) 77
- ISC-D 77
- ISC-M 77

K

- key-controlled storage protection 15

L

- large send 70
- Layer 2 (Link Layer) 70
- Layer 3 virtual MAC 70
- LCSS (logical channel subsystem) 31
- Licensed Internal Code (LIC) 111
- links
 - InfiniBand (IFB) 78
 - Integrated Cluster Bus-4 (ICB-4) 78
 - internal coupling (IC) 80
 - ISC-3 77
- Linux on System z 30
 - VLAN support 71
- logical channels subsystem (LCSS) 31
- logical partition (LPAR)
 - increased 4
- logically partitioned operating mode 22
- LPAR
 - definition 42
- LPAR (logical partition)
 - clustering 89
- LPAR mode 22

M

- machine information 136
- machine-check interruptions 58
- maximums
 - channel, ports, adapters 31
- MBA (Memory Bus Adapter) fanout cards 17
- memory
 - central storage 15
 - expanded storage 15
 - rules 16
- Memory Bus Adapter (MBA)
 - failure containment 122
- Memory Bus Adapter (MBA) fanout cards 17

- memory cards
 - characteristics 16
- memory scrubbing 124
- MIDAW (Modified Indirect Data Address Word)
 - facility 48
- MIF (Multiple Image Facility) 40
- modes
 - burst 55
 - byte multiplexer 55
 - channel multiplexing 55
 - LPAR 22
 - operation 57
 - peer 77
- Modified Indirect Data Address Word (MIDAW)
 - facility 48
- Modify Subchannel (MSCH) instruction 57
- multipath IPL 48
- Multiple Image Facility (MIF) 40

N

- network concentrator 71
- Network Time Protocol (NTP) 85
- network traffic analyzer (NTA)
 - HiperSockets 71, 131
 - OSA-Express 69, 131
- nonsynchronous operation 62
- NTP (Network Time Protocol) 85

O

- On/Off CoD (On/Off Capacity on Demand) 138
- online information xviii
- operating system messages 110
- operation
 - block multiplexer 55
 - byte multiplexer 55
 - ESA/390 57
 - nonsynchronous 62
 - retry 58
 - synchronous 61
- OSA LAN idle timer 68
- OSA-Express network traffic analyzer 69, 131
- OSA-Express2 67
- OSA-Express3 65
- OSA/ETR card 17
- OSA/SF (OSA/Support Facility) 65
- OSA/Support Facility (OSA/SF) 65
- OSC (oscillator) card 12, 17
- oscillator (OSC) card 12, 17

P

- Parallel Sysplex 75
 - coupling link connectivity 76
- partial memory restart 125
- patch apply 81
- PCHID
 - assignments 37, 38
 - CPC drawer 38
 - I/O drawer 38

PCHID (continued)

- report, sample 39
- peer channels 33
- peer mode 77
- performance
 - device 55
 - ESCON channel 61
 - ESCON characteristics 63
 - I/O 55
 - system 7
- permanent upgrades 137
- planning
 - channel subsystem 37
- ports
 - maximums 31
- POS 98
- positioning
 - FCP 52
- positions
 - drawers 12
- power estimation tool 136
- power sequence control (PSC) 26
- power supply 21
- PPS (pulse per second) PPS 85
- PR/SM (Processor Resource/Systems Manager) 88
- PR/SM LPAR
 - CPU management (Clustering) 89
 - Sysplex Timer offset 24
- Preferred Time Server 126
- priority queuing 89
- problem analysis and reporting 110
- Processor Resource/Systems Manager (PR/SM) 22, 88
- processor unit (PU) 12, 23
 - sparing 122
- programming
 - compatibility 27
 - support 27
- protocol
 - I/O
 - data streaming 55
 - interlocked 55
 - interlocked 55
- PSC (power sequence control) 26
- PSC24V card 20
- PU (processor unit) 12, 23
- publications
 - Cryptographic xvii
 - miscellaneous xvii
 - OSA-Express xvi
 - parallel sysplex xvi
 - prerequisite xv
 - related xv
- pulse per second (PPS) signal 85
- purge path extended 48

Q

- Queued Direct I/O Diagnostic Synchronization (QDIOSYNC) 69

R

- RAS (Reliability, Availability, Serviceability)
 - availability 121
 - reliability 121
 - serviceability 132
- reconfigurable channels 23
- reconnection, dynamic 61
- related
 - publications xv
 - websites xvii
- remote automated operations 119
- remote key loading 98
- remote operations
 - manual 118
 - using a Hardware Management Console 119
 - using a web browser 119
 - overview 117
- Remote Support Facility (RSF) 112
- reports
 - cabling 102
 - PCHID 39
- Reset Channel Path (RCHP) instruction 57
- reset, system 58
- Resource Link 7
- Resume Subchannel (RSCH) instruction 57
- retry
 - operations 58
- RMF monitoring 98
- RSF (Remote Support Facility) 112
- rules
 - memory 16

S

- safety notices xiii
- sample
 - PCHID report 39
- SAP (System Assist Processor) 14
- scheduled operations 111, 112
- SCM (single chip module) 12
- SCSI (Small Computer Systems Interface) 49
- SE (Support Element) 20, 105
- security 116
 - CPACF 93
 - cryptographic accelerator 94
 - cryptographic coprocessor 94
- sending comments xviii
- Server Time Protocol (STP)
 - description 25, 83
- service required state 109
- Set Channel Monitor (SCHM) instruction 57
- shared channels 23
- SIE (start interpretive execution) instruction 28
- single chipmModule (SCM) 12
- Small Computer Systems Interface (SCSI) 49
- software support 27
 - Linux on System z 30
 - TPF 30
 - z/OS 28
 - z/OS.e 28

- software support (*continued*)
 - z/VM 29
 - z/VSE 29
- spanned channels 23, 41
- Start Subchannel (SSCH) instruction 57
- status reporting 108
- storage
 - central 15, 23
 - expanded 15, 23
 - z/Architecture 15
- Store Report Word (STCRW) instruction 57
- Store Subchannel (STSCH) instruction 57
- subchannel
 - connectivity 34
- support
 - broadcast 71
 - operating systems 27
 - software 28
- Support Element (SE) 20, 105
- synchronous operation 61
- sysplex functions
 - parallel 75
- Sysplex Timer attachment 17
- Sysplex Timer offset 24
- system
 - configurations 11
 - I/O reset 58
- System Assist Processor (SAP) 14
- system power supply 21
- System z10 Application Assist Processor (zAAP) 14
- system-managed coupling facility structure duplexing (CF duplexing) 86

T

- Test Pending Interruption (TPI) instruction 57
- Test Subchannel (TSCH) instruction 57
- time-out functions
 - channel 59
- TKE (Trusted Key Entry) 96
- tools
 - CHPID mapping 40
- TPF 30
- Trusted Key Entry (TKE) 96

U

- unsupported features 9
- upgrade progression 9
- upgrades
 - nondisruptive 140
 - permanent 137

V

- virtual RETAIN 111

W

- websites xvii
- workload manager 90
- WWPN prediction tool 136

Z

- z/Architecture 8
- z/OS 28
- z/OS.e 28
- z/TPF 30
- z/VM 29
 - dynamic link aggregation 69
- z/VSE 29
- zAAP (System z10 Application Assist Processor) 2, 14
- zIIP (System z10 Integrated Information Processor) 2



Level 04a, October 2010
Printed in USA

SA22-1085-04

